

INTERNATIONAL COOPERATION TREATY

PCT

NOTIFICATION OF ELECTION

(PCT Rule 61.2)

From the INTERNATIONAL BUREAU

To:

Assistant Commissioner for Patents
 United States Patent and Trademark
 Office
 Box PCT
 Washington, D.C.20231
 ETATS-UNIS D'AMERIQUE

in its capacity as elected Office

Date of mailing (day/month/year) 25 April 2000 (25.04.00)	
International application No. PCT/GB99/03229	Applicant's or agent's file reference 5279399/EMR
International filing date (day/month/year) 29 September 1999 (29.09.99)	Priority date (day/month/year) 29 September 1998 (29.09.98)
Applicant SLATER, Mel	

1. The designated Office is hereby notified of its election made:

☒ in the demand filed with the International Preliminary Examining Authority on:

10 March 2000 (10.03.00)

☐ in a notice effecting later election filed with the International Bureau on:2. The election ☒ was☐ was not

made before the expiration of 19 months from the priority date or, where Rule 32 applies, within the time limit under Rule 32.2(b).

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No.: (41-22) 740.14.35	Authorized officer Olivia RANAIVOJAONA Telephone No.: (41-22) 338.83.38
---	---

PCT

REC'D 22 DEC 2000

WIPO

PCT

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

(PCT Article 36 and Rule 70)

Applicant's or agent's file reference EMR/ACS/5279399		See Notification of Transmittal of International Preliminary Examination Report (Form PCT/IPEA/416) FOR FURTHER ACTION
International application No. PCT/GB99/03229	International filing date (day/month/year) 29/09/1999	Priority date (day/month/year) 29/09/1998
International Patent Classification (IPC) or national classification and IPC G06T15/50		
Applicant UNIVERSITY COLLEGE LONDON et al.		

1. This international preliminary examination report has been prepared by this International Preliminary Examining Authority and is transmitted to the applicant according to Article 36.



2. This REPORT consists of a total of 7 sheets, including this cover sheet.

- ☐ This report is also accompanied by ANNEXES, i.e. sheets of the description, claims and/or drawings which have been amended and are the basis for this report and/or sheets containing rectifications made before this Authority (see Rule 70.16 and Section 607 of the Administrative Instructions under the PCT).

These annexes consist of a total of sheets.

3. This report contains indications relating to the following items:

- I ☒ Basis of the report
- II ☐ Priority
- III ☐ Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- IV ☐ Lack of unity of invention
- V ☒ Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- VI ☐ Certain documents cited
- VII ☒ Certain defects in the international application
- VIII ☒ Certain observations on the international application

Date of submission of the demand 10/03/2000	Date of completion of this report 20.12.2000
Name and mailing address of the international preliminary examining authority:  European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Tx: 523656 epmu d Fax: +49 89 2399 - 4465	Authorized officer Meinl, W Telephone No. +49 89 2399 2532 

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

International application No. PCT/GB99/03229

I. Basis of the report

1. This report has been drawn on the basis of *(substitute sheets which have been furnished to the receiving Office in response to an invitation under Article 14 are referred to in this report as "originally filed" and are not annexed to the report since they do not contain amendments (Rules 70.16 and 70.17).):*

Description, pages:

1-80 as originally filed

Claims, No.:

1-57 as originally filed

Drawings, sheets:

1-32 as originally filed

2. With regard to the **language**, all the elements marked above were available or furnished to this Authority in the language in which the international application was filed, unless otherwise indicated under this item.

These elements were available or furnished to this Authority in the following language: , which is:

- ☐ the language of a translation furnished for the purposes of the international search (under Rule 23.1(b)).
- ☐ the language of publication of the international application (under Rule 48.3(b)).
- ☐ the language of a translation furnished for the purposes of international preliminary examination (under Rule 55.2 and/or 55.3).

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, the international preliminary examination was carried out on the basis of the sequence listing:

- ☐ contained in the international application in written form.
- ☐ filed together with the international application in computer readable form.
- ☐ furnished subsequently to this Authority in written form.
- ☐ furnished subsequently to this Authority in computer readable form.
- ☐ The statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.
- ☐ The statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished.

4. The amendments have resulted in the cancellation of:

- ☐ the description, pages:
- ☐ the claims, Nos.:

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

International application No. PCT/GB99/03229

☐ the drawings, sheets:

5. ☐ This report has been established as if (some of) the amendments had not been made, since they have been considered to go beyond the disclosure as filed (Rule 70.2(c)):

(Any replacement sheet containing such amendments must be referred to under item 1 and annexed to this report.)

6. Additional observations, if necessary:

V. Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-57
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-57
Industrial applicability (IA)	Yes:	Claims	1-57
	No:	Claims	

2. Citations and explanations
see separate sheet

VII. Certain defects in the international application

The following defects in the form or contents of the international application have been noted:
see separate sheet

VIII. Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:
see separate sheet

Re. Sections V, VII, VIII (Novelty, Inventive step, Defects)

1. The following documents are referred to in this report:-

- D1:** LEVOY M et al: 'Light Field Rendering', COMPUTER GRAPHICS PROCEEDINGS (SIGGRAPH), US, NEW YORK, NY: ACM, 4 August 1996 (1996-08-04), page 31-42, cited in the application
- D2:** GORTLER S J et al: 'THE LUMIGRAPH', COMPUTER GRAPHICS PROCEEDINGS (SIGGRAPH), US, NEW YORK, NY: ACM, 4 August 1996 (1996-08-04), page 43-54, cited in the application
- D3:** EP-B-0 795 164
- D4:** US-A-5,488,700
- D5:** NEUMANN L et al: 'Radiosity and Hybrid Methods', ACM TRANSACTIONS ON GRAPHICS, US, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, vol. 14, no. 3, July 1995 (1995-07), page 233-265, ISSN: 0730-0301
- D6:** EP-A-0 610 004
- D7:** US-A-5,313,568
- D8:** NG A: 'Assessment of five radiosity acceleration techniques', COMPUTERS AND GRAPHICS, GB, PERGAMON PRESS LTD. OXFORD, vol. 19, no. 5, September 1995 (1995-09) - October 1995 (1995-10), page 727-738
- D9:** NEUMANN L et al: 'Radiosity with Well Distributed Ray Sets', THE INTERNATIONAL JOURNAL OF THE EUROGRAPHICS ASSOCIATION, COMPUTER GRAPHICS FORUM, vol. 16, no. 3, 4 - 8 September 1997, pages C261-C269, Budapest

2. The various definitions of the invention given in independent device claims 1, 14 and 24 are such that the claims as a whole are not clear and concise, contrary to Article 6 PCT. The same applies to the independent method claims 28, 40 and 49. The claims should be recast to include only the minimum necessary number of independent claims in any one category, with dependent claims as appropriate, Rule 6.4 PCT.

While it is true that the apparatus claims 1, 14 and 24 have the core features in common and differ in substance only in their opening phrases, the necessity of three independent apparatus claims for virtually the same thing cannot be seen. For instance, no real difference is apparent between an "apparatus for processing data" (claim 1) which inherently also "generates data" and an "apparatus for generating data" (claim 14).

3. Brief review of the prior art.

The documents **D1**, **D2** concern image-based "light field" rendering.

The document **D3** concerns ray tracing including an algorithm that reduces memory accesses when determining intersections of the eye rays with polygons.

The documents **D4 - D8** concern hybrid methods including radiosity and ray tracing, e.g. with ray tracing for finding the form factors as in D7 or D8.

The document **D9** concerns a radiosity method with Monte Carlo techniques.

4. Novelty.

- 4.1 Claim 1 is so broad that it is anticipated by any standard ray-tracing and radiosity method for determining global illumination. As should be acknowledged by the applicant who is expert in the field of computer graphics, each of those methods includes defining 3D objects, light sources and calculating received "energy" i.e. light intensities or radiosity by taking into account intersections of the light rays with the 3D objects, equally as specified in claim 1.

A good summary on ray tracing can be found in **D4**, col.1-3, and on radiosity in **D8**, pages 727-729; see also the text book of **Foley, van Dam et al**, 2nd ed., chapters 16.12 and 16.13.

It thus would appear that the subject-matter of claim 1 lacks novelty over each of the documents **D3 - D9** and also over Foley, van Dam (Art.33(2) PCT).

Regarding the question of novelty over D3 - D9 it is irrelevant that these documents are not concerned with "light fields" as long as the claimed features have antecedents in these documents which evidently is the case.

- 4.2 Regarding claims 2 and 3 it is not clear in as far the claims contain any limitation with respect to claim 1, and in as far they differ from standard ray tracing or radiosity methods. For instance, a "viewing plane within said [3D] environment" as recited in claim 3 is a usual feature in VR software that allows walking through an environment. The claims thus also seem to lack novelty.

4.3 Claim 4.

Determination of "angles of incidence" with the viewing plane is a feature that seems inherent in ray tracing, see for instance Foley, van Dam, 2nd. ed.,

Fig.15.55. The geometry in Fig. 15.55 regarding the center of projection (eye point) and the viewing plane inherently gives the angles of incidence between the light rays and the viewing plane. Thus, ray tracing alone anticipates the subject-matter of claim 4.

5. Inventive step.

Even were claim 1 interpreted to clearly address "light fields" the claim would appear to lack an inventive step with respect to each of D1 and D2.

D1 and D2 are concerned with creation of a light field from 2D image data which is generated from either rendering or from digitising. In contrast, the claimed invention defines objects in a 3D space (lines 5, 6 and 9, 10 of claim 1), thus establishing novelty over D1 and D2.

It however seems that the remaining features of claim 1 (apart from the distinguishing features) merely relate to "light fields" as in D1 or D2, and that claim 1 merely expresses the wish to process 3D scenes with the recently developed light field technique. Differently stated, claim 1 merely claims the extension of the known image-based light fields to 3D light fields without giving further details of the implementation. The wishful formulation of an extension from 2D to 3D is clearly obvious in computer graphics.

The dependent claims not appear to contain any clear feature that would go beyond the "light fields" of D1 or D2.

Further remarks on clarity and "light fields" see point 7. below.

6. The above objections apply to the other independent claims accordingly.

7. Clarity.

It seems that the "light fields" are an essential feature of the claimed invention (see pp. 21, 22 of description) and it is noted that neither the "light fields" nor any features clearly related thereto are recited in any of the independent claims, contrary to Art.6 PCT.

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT - SEPARATE SHEET**

International application No. PCT/GB99/03229

When amending, it would thus appear necessary that, broadly speaking, the single independent apparatus claim be limited to "light fields" as this concept is called in D1. However, since it appears that the term "light field" was not a very common term at the filing date of the application (apart from the single reference D1), the term should be duly specified by technical features.

PCT


WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06T 15/50		A1	(11) International Publication Number: WO 00/19379
			(43) International Publication Date: 6 April 2000 (06.04.00)
(21) International Application Number: PCT/GB99/03229		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 29 September 1999 (29.09.99)		Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	
(30) Priority Data: 9821139.4 29 September 1998 (29.09.98) GB 9915838.8 6 July 1999 (06.07.99) GB			
(71) Applicant (for all designated States except US): UNIVERSITY COLLEGE LONDON [GB/GB]; Gower Street, London WC1E 6BT (GB).			
(72) Inventor; and (75) Inventor/Applicant (for US only): SLATER, Mel [GB/GB]; 2 Arlington Gardens, Ilford, Essex IG1 3HH (GB).			
(74) Agents: BERESFORD, Keith, Denis, Lewis et al.; Beresford & Co., 2-5 Warwick Court, High Holborn, London WC1R 5DJ (GB).			

(54) Title: ENERGY PROPAGATION MODELLING APPARATUS



(57) Abstract

A simulation apparatus (2, 70) is operable to simulate propagation of light within a scene. The apparatus defines an environment containing a plurality of discrete paths (72) in a plurality of directions, along which energy is propagated in the environment. The apparatus determines the positions of objects and light sources in a scene, and records points (t) on the paths, at each of which points there is an intersection of an object or light source within the scene with the path (72). The apparatus identifies propagation of light source from the light sources, along the paths, and to further paths in accordance with the intersections. A corresponding method is also provided. Distribution of propagation of other energy, such as sound or heat, can be simulated in addition to or as an alternative to light propagation simulation.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

ENERGY PROPAGATION MODELLING APPARATUS

The present invention is concerned with apparatus for modelling the propagation of energy. In particular, the present invention is concerned with the modelling of propagation which can be represented by ray paths, and the manner in which energy passes along ray paths while interacting with objects in a modelled scene. The present invention has applications for example in the modelling of light in a scene, to provide a facility for production of near photo-realistic images of a globally illuminated scene or to provide for measurement of light or other energy distribution within a scene.

15 In many computer games, dynamic imagery is of real importance, and so real-time production of dynamic images is implemented to the detriment of photo-realism. On the other hand, many applications such as architectural design do not require substantial dynamic changes to an image in real-time, whereas photo-realism is of real importance.

However, even in a photo-realistic model, it is preferable that a user can change the viewpoint of a computer modelled scene, with limited re-calculation. An ideal technique, for carrying this out without

significant re-calculation, is known as a "view-independent" technique. View-independence in this sense means that data pertaining to the scene does not require new lighting calculations on variation of the point of view of the scene.

In order to model a scene in a photo-realistic manner, the true reflective and transmissive properties of light incident on objects in the scene must be approximated. In particular, the reflection and transmission of light must be modelled so as to approximate to the behaviour of light in real life. In that way, the observer of the modelled scene will reach the correct conclusions concerning the nature and appearance of objects within the scene.

Light within a scene has in the past been modelled by several techniques, most of which can be categorised as either "ray tracing" techniques or "radiosity" techniques.

Ray tracing is described in detail in "An Improved Illumination Model for Shaded Display" by T. Whitted, published in COMM.AMC 23(6), 343-349, June 1980. Ray tracing assumes that the observer is viewing an image projected through a pin-hole camera, on to an image

plane. The image is of a scene illuminated by point light sources. For ease of calculation, the medium through which light is transmitted is usually assumed to be non-participating, but light attenuation effects can be modelled as described in "Computer Graphics, Principles and Practice" by Foley, van Dam, Feiner, Hughes, 2nd ed. (hereinafter referred to as "Foley") pages 727 and 728.

10 Rays are traced from the pin-hole or centre of projection (COP) through the image plane and into the scene, bouncing from object to object using the laws of specular reflection (angle of reflection equals angle of incidence, incident and reflected rays are co-planar),
15 until the ray exits the scene or it is established that further reflections will make negligible contribution to the image to be displayed on the computer screen because the intensity of further reflected rays is negligible. In practice, when considering a scene including
20 transparent or translucent objects, transmission of light is considered in the same way as reflection, using well-known physical laws.

Ray tracing is entirely view-dependent. That means that
25 if a different viewpoint is selected, the transmission of light towards the viewpoint from the scene must be

entirely re-calculated. Moreover, it is not suitable for modelling diffuse reflection, where a particular incident ray results in a plurality of reflected ray directions, and a particular reflected ray can be the result of a plurality of incident rays. It is therefore impossible to represent diffuse reflection realistically using ray tracing.

Radiosity techniques are discussed in the paper "Modelling the interaction of light between diffuse surfaces" by C M Goral et al, published in Computer Graphics (SIGGRAPH 84), 213-222. Radiosity techniques are suitable for scenes where all materials are ideal diffuse reflectors. However, radiosity is not suitable for consideration of specular reflection. Radiosity is defined as the rate at which light energy leaves a surface.

Almost all radiosity methods rely on a scene to be modelled being described in terms of polygons. Each polygon can be further sub-divided into small patches. The patches may in turn be sub-divided into surface elements. Each patch (or surface element as the case may be) is considered in turn, and the radiosity at that patch is calculated from the radiosity of patches from which light can be transmitted to the patch in question,

taking into account their relative position and angular orientations. It should be appreciated that radiosity techniques cannot be used for point light sources, since a light source must be capable of description in terms of polygons defining the geometry of the light source.

Once the radiosity at each patch has been calculated, Gouraud shading can be employed to produce a smooth image. Gouraud shading is described in Foley, pp 736 and 737. Since patches of finite area are considered instead of points on a surface, radiosity will not result in accurate consideration of specular reflection. However, the lack of specular reflective effects means that the lighting of a model considered through radiosity techniques is independent of the position of the viewpoint of the model. That is, recalculation of the whole image need not take place in the event that the position of the viewpoint changes.

As a result, a radiosity model is suitable to be viewed from a dynamic viewpoint in real time. This is known as "real-time walk-through". This is highly suitable for CAD applications such as for architectural models, but the loss of specular reflection results in a loss in photo-realism. It should be appreciated that significant recalculation will be necessary at every change of view

due to the need to determine visible surfaces and to apply interpolative shading.

Finally, a rendering technique proposed on the one-hand
5 by Levoy and Hanrahan in a paper published in Computer Graphics (SIGGRAPH), Annual Conference Series (1996), 31-42, and on the other-hand by Gortler et al in Computer Graphics (SIGGRAPH), Annual Conference Series (1996), 43-52, does not fall within the scope of either ray
10 tracing or radiosity methods.

In both disclosures, a system implementing a "light field" or "lumigraph" is described. The light field is constructed of an ordered plurality of rays constructed
15 between two parallel planes. Other light fields can be constructed using different arrangements of construction planes. The rays are arranged between discrete points on the planes, and can be described in terms of the points on the two planes between which they extend. A scene
20 consisting of one or more objects can be rendered into the light field, by considering views of the scene at each point on the two planes. Those views can either be constructed by one of the techniques identified above, or by importing real views of a scene produced on digital
25 cameras. The radiance propagated along each ray from a particular point constructing the light field can be

identified from the view from that point constructed or imported by the system.

Once the radiances associated with rays from each point
5 have been identified, views of the scene from a point within the light field can be constructed by considering the radiance of rays passing near that point and interpolating an image from the radiance values obtained.

10 Real images of real objects can be rendered into the light field, so that images from positions other than the positions from which images were obtained by a camera can also be seen. These real images can be generated by means of a digital camera, taking digital photographs
15 from a plurality of different positions about the scene to be rendered into the light field.

However, the light field technique described above has a significant disadvantage, in that views cannot be
20 constructed from every position in the light field. If a scene consists of a plurality of objects, or at least one object which has at least one concave surface, then a volume can be defined, between a surface bounding a minimum volume about the scene which includes no concave
25 surfaces (the "convex hull" of the scene) and the actual surface(s) defining the object(s) of the scene. Several

rays within this volume are bounded at both ends by surfaces of objects in the scene. Therefore, radiances along those rays cannot be determined from consideration of real images at points on the two constructing planes.

5 As a result, discontinuities exist which prevent construction of images from a sub-set of all viewpoints within the volume.

It is an object of the present invention to provide an
10 alternative technique for modelling energy propagation, such as the propagation of light through a scene.

The invention provides simulation apparatus which defines discrete paths within an environment within which energy
15 propagation is to be modelled, and which locates points along those paths which denote interactions of objects within the environment with the paths.

One aspect of the present invention provides a technique
20 which allows specular and diffuse reflection to be represented realistically. A scene, and light propagation within the scene, can be represented by the present invention such that an image of the scene from a particular viewpoint can be generated in substantially
25 constant time. The constant time is not dependent on the complexity of the scene to be represented and, with the

provision of sufficient computer power, is capable of being made significantly small that real time changes of viewpoint can be represented.

- 5 Further features and advantages of the invention will become apparent from the following description of apparatus and a method in accordance with a preferred and specific embodiment of the invention, with reference to the accompanying drawings in which:

10

Figure 1 is a schematic diagram of image processing apparatus in accordance with a specific embodiment of the invention;

- 15 Figure 2A is a perspective view of a scene to be modelled by the image processing apparatus of Figure 1, including illustration of ray directions encountered at a first viewpoint;

- 20 Figure 2B is a perspective view of the scene illustrated in Figure 2A, including illustration of ray directions encountered at a second viewpoint;

- Figure 3A is a perspective view of a parallel sub-field
25 in accordance with a specific embodiment of the present invention;

Figure 3B is a perspective view of a further parallel sub-field in accordance with a specific embodiment of the present invention;

- 5 Figure 3C is a perspective view of a still further parallel sub-field in accordance with a specific embodiment of the present invention;

Figure 4 is a perspective view of the scene illustrated
10 in Figures 2A and 2B, including illustration of a selected plane within the scene to be modelled and a selected ray within that plane in accordance with the specific embodiment of the present invention;

- 15 Figure 5 is an elevation of the construction plane illustrated in Figure 4 in a direction normal to the construction plane, in accordance with the specific embodiment of the present invention;

- 20 Figure 6 is a schematic diagram illustrating intersections of objects with the light ray illustrated in Figure 4, throughout its length;

Figure 7 is a schematic diagram showing an extract of the
25 length of the light ray, between the side walls of the scene illustrated in Figure 4;

11

Figure 8 is a schematic diagram showing the structure of interval data in accordance with the specific embodiment of the present invention;

5 Figure 9 is a schematic diagram of a data structure for the interval data illustrated in Figure 8;

Figure 10 is an elevation similar to that illustrated in Figure 5, including a representation of a virtual eye for
10 viewing the scene;

Figure 11 is a schematic diagram showing the structure of the image processing apparatus illustrated in Figure 1;

15 Figure 12 is a schematic diagram showing the structure of the L F computation unit illustrated in Figure 11;

Figure 13 is a schematic diagram showing the internal structure of the viewer illustrated in Figure 11;
20

Figure 14 is a flow diagram illustrating the procedure performed by the image processing apparatus in use;

Figures 15 to 20 are flow diagrams illustrating sub-procedures called by the procedure illustrated in Figure
25 14;

Figure 21 is a flow diagram illustrating a procedure according to which the viewer operates in use;

Figure 22 is a flow diagram illustrating a procedure by
5 which the image processing apparatus is operative to delete an object from a scene;

Figure 23 is a flow diagram illustrating a procedure by
which the image processing apparatus is operative to add
10 an object to a scene;

Figure 24 is a schematic diagram of an object in a light field demonstrating the potential for computational complexity in considering diffuse reflection;
15

Figure 25 is a schematic diagram showing a portion of the surface of the object illustrated in Figure 24, illustrating a gathering step of a method of reducing computational complexity;
20

Figure 26 is a diagram of the portion of the object surface illustrated in Figure 25, demonstrating a "shooting" step of the method of reducing computational complexity;
25

Figure 27 is a view of the object illustrated in Figure

24, demonstrating a further method of reducing the computational complexity of considering diffuse reflection;

5 Figure 28 is a schematic diagram of image processing apparatus for defining an environment in virtual reality, in accordance with the specific embodiment of the invention;

10 Figure 29 is a side view of a head mounted display for use in the apparatus illustrated in Figure 28; and

Figure 30 is a view of the display apparatus of the head mounted display illustrated in Figure 29.

15

Figure 1 is a block diagram showing the general arrangement of an image processing apparatus according to an embodiment of the invention. In the apparatus, there is provided a computer 2, which comprises a central processing unit (CPU) 4 connected to a memory 6 operable to store a program defining the sequence of operations of the CPU 4 and to store object and image data used in calculations by the CPU 4.

20

25 An input device 8 is coupled to an input port (not shown) of the CPU 4. The input device 8 may comprise, for

example, a keyboard and/or a position sensitive input device such as a mouse, tracker ball, or a digitizer tablet and stylus etc.

5 A frame buffer 10 is also coupled to the CPU 4, the frame buffer 10 comprising a memory unit (not shown) arranged to store image data relating to at least one image, for example by providing one (or several) memory location(s) per pixel of the image. The value(s) stored in the frame
10 buffer 10 for each pixel defines the colour or intensity of that pixel in the image.

In the present embodiment of the invention, an image is represented as a two-dimensional array of pixels, which
15 can conveniently be described in terms of Cartesian co-ordinates. The position of the given pixel can be described by a pair of x, y co-ordinates. The frame buffer 10 has sufficient memory capacity to store at least one image. If the image has a resolution of 1000
20 by 1000 pixels, the frame buffer 10 should include 10^6 pixel locations, each location being addressable directly or indirectly in terms of pixel co-ordinates x, y.

A video display unit (VDU) 12 is coupled to the frame
25 buffer 10. The VDU 12 is operable to display the image stored in the frame buffer 10 in a conventional manner.

For instance, if the VDU 12 displays images in a raster scanning manner, the x co-ordinate of a pixel maps to the distance along a line of the scanned display, and the y co-ordinate of the pixel maps to the number of the line.

5

Also coupled to the frame buffer 10 is a video tape recorder (VTR) 14 adapted to receive a video tape 15. Another image recording device, such as a paper printer, a 35mm film recorder or a recordable compact disc could be provided in addition or in the alternative.

A mass storage device 16, such as a hard-disk drive, is coupled to the memory 6. The mass storage device 16 has a high data storage capacity, suitable for storing data to which instant access is not required. Moreover, a disk drive 18, operable to accept removable data storage media such as a floppy disk 20 is coupled to the memory 6. The disk drive 18 is operable to transfer data stored on the floppy disk 20 to the memory 6.

20

The apparatus is programmable through program instructions introduced on a storage medium, such as a floppy disk 20. Alternatively, an optical disk reader could be provided to receive an optical disk, such as in CD-ROM or DVD format. Further, the apparatus could include a modem, to receive a signal bearing program

25

instructions. The program instructions could also be introduced by typing at the keyboard of the input device 8.

5 Figure 2A illustrates a scene 22 in respect of which it would be desirable to model light propagation using the apparatus of the embodiment of the invention described herein. The scene 22 consists of a room bounded by left and right side walls 24, 26, back and front walls 28, 30,
10 a floor 32 and a ceiling 34. The front wall 30 and the ceiling 34 are shown as transparent in Figure 2A for reasons of clarity. The walls, floor and ceiling together define a room which, in this embodiment, is substantially cubic in shape.

15

A window 36 is set in the left-hand side wall 24, and a rectangular mirror 38 is mounted on the back wall 28. A free-standing lamp 40 stands on the floor 32, in the corner defined by the left-hand side wall 24 and the back
20 wall 28. A chair 42 is situated substantially in the centre of the floor 32. The chair 42 is oriented in the room such that its back is substantially parallel with a line constructed between the back left corner of the room and the front right corner of the room.

25

In Figure 2A, an eye 44 is placed at a viewpoint at the

front of the scene 22. As is well-known from conventional optics, light enters the eye 44 and is refracted by various elements of the eye 44, to allow a focused image to be formed on the retina.

5

In the scene 22, rays of light travel in an infinite number of directions. However, only those rays which enter the pupil of the eye 44 are experienced by the observer. For example, a ray 46 is shown travelling from the window 36 (transmitting light from a theoretically infinitely distant light source), directly to the eye 44. Another ray 48 is shown travelling directly from the lamp 40 to the eye 44.

15 However, a substantial amount of light will not reach the eye 44 directly from light sources, but will be reflected off objects within the scene 22. For example, light radiated from the lamp 40 will hit the left-hand wall 24, the floor 32 and the right-hand wall 26, and be reflected towards the eye 44. Rays 50a, 50b, 50c respectively illustrate these three circumstances. It will be understood that light radiated from the lamp will also be reflected towards the eye from the back wall 28, the front wall 32 and the ceiling 34, but rays corresponding to those paths are omitted from Figure 2A for reasons of clarity.

25

Moreover, rays 52a, 52b in Figure 2A represent light travelling from the window 36 to the eye 44 via a reflection off the back of the chair 42. Rays 54a, 54b represent light travelling from the lamp 40 to the eye 44 via a reflection off the seat of the chair 42.

The walls, floor and ceiling 24-34 and the chair 42 have diffusely reflecting surfaces. Consequently, a ray of light incident on one of those objects will reflect in all directions forming an acute angle with the surface normal at the point of incidence. Unit vectors representing those directions define a hemispherical volume of unit radius, with centre at the point of incidence, and bounded by a plane tangential to the surface of the object at the point of incidence.

As a result, an observer of light reflected at one of the objects with diffusely reflective surfaces will see a representation of the object in question. The observer will not see a sharp reflected image of the ultimate source of the reflected light. Also, some of the light incident on each of these objects will be absorbed thereby. The rate of absorption is normally frequency dependent, and the absorption is associated with interpretation by the observer of colour of the object in question.

In contrast, the ray 56 illustrated in Figure 2A, which is firstly diffusely reflected from the back of the chair 42 towards the mirror 38, is then specularly reflected from the mirror towards the eye 44. Specular reflection is that which behaves in accordance with the ideal laws of reflection, as described in the introduction above. By observing other rays which similarly diffusely reflect from the various objects in the scene to the mirror and then to the eye, a reflected image of the scene is formed in the eye.

Referring now to Figure 2B, the same scene 22 is illustrated, with the eye of the observer in a new position. The eye is now referenced with the numeral 44'. The objects within the scene 22 described above are given the same reference numerals as before, in view of the fact that they are unchanged. However, the scene 22 is now being viewed from a different position, and so different rays of light will enter the eye 44'. For example, a ray 58 travels directly from the window 36 to the eye 44', and a ray 60 travels directly from the lamp 40 to the eye 44'. Furthermore, rays 62a, 62b travel from the lamp 40 via diffuse reflection off the chair 42 to enter the eye 44'. A ray 64 will travel from the lamp 40, diffusely reflect from the back wall 28, and enter the eye 44'. A ray 66 from the lamp 40 will specularly

reflect from the mirror 38 and enter the eye 44'. That ray 66 is representative of rays from the lamp which reflect from the mirror 38, which will contribute to the construction of a reflected image of the lamp within the eye.

From the foregoing description of the two views of the same scene, it will be seen that in order to observe the same scene from two different views, one must obtain information concerning the radiance of light associated with two different sets of rays.

In the real world, light exists independently of a viewer, and behaves in accordance with physical laws such as concerning reflection, refraction and diffraction. Where a view of a scene is observed in the real world, those rays which enter the aperture of a viewer are taken into consideration, and focused to produce an image. The preferred embodiment of the present invention aims to emulate that situation, by modelling a complete pattern of light behaviour within a scene, without reference to a particular viewpoint.

This objective is not achievable in practice, because any scene contains an infinite number of potential light propagation directions. Therefore, the embodiment models

a light field consisting of a finite number of potential carriers of light radiance, referred to hereinafter as "rays".

5 Thereafter, objects can be rendered into the light field by considering the intersections of the rays in the light field with the objects. Once all objects have been rendered into the light field, light can be propagated along the rays from objects designated as light emitters,
10 and interactions such as reflections, refractions and defractions can be computed at the intersections of the rays with the objects.

The embodiment is significantly different from the
15 techniques described by Levoy and Hanrahan and Gortler et al, both of which are directed to the construction of a 3D representation of a scene from 2D images. In contrast, the preferred embodiment described herein is concerned primarily with the behaviour of light within
20 the light field, and considers the characteristics of the objects rendered into the light field only when calculating interactions at intersections.

Series of 2D images need not be stored with the present
25 invention; a final image at a particular viewpoint in the light field can be constructed in a straightforward

manner.

Various different light fields could be defined, but a preferred light field consists of a plurality of parallel sub-fields. Each parallel sub-field consists of a light region, of cuboid shape, between two opposite sides of which extend a plurality of parallel rays. The parallel rays are arranged in a grid to enable straightforward referencing thereof. The centre of each parallel sub-field is defined as its origin, and the parallel sub-fields are overlaid so that their origins coincide. The parallel sub-fields are oriented in different directions in relation to each other, so that a light field is constructed which includes rays offset across a region, and oriented in a variety of directions in three-dimensional space.

With reference now to Figure 3A, a parallel sub-field comprises a cubic volume in x, y, z Cartesian space defined by:

$$-1 \leq x \leq 1$$

$$-1 \leq y \leq 1$$

$$-1 \leq z \leq 1$$

25

A rectangular grid is imposed on the x, y plane

comprising n sub-divisions of equal width in each of the x and y directions. The width of each sub-division in each of the x and y directions is $2/n$. The value of n is defined by reference to the level of resolution of final
 5 images generated in the apparatus which would be acceptable, and the amount of data storage of the apparatus, in particular the mass storage unit 16 thereof. In this embodiment, n is set at 200.

10 Discrete co-ordinates (i, j) are assigned to the centre of these sub-divisions, imposing a discrete co-ordinate grid on the x, y plane within the parallel sub-field 70. The grid is referenced by co-ordinates (i, j) . The real co-ordinate on the x, y plane corresponding to (i, j) is
 15 therefore:

$$(x_i, y_j) = \left(\frac{2i+1}{n} - 1, \frac{2j+1}{n} - 1 \right); i=0, 1, \dots, n-1; j=0, 1, \dots, n-1$$

Each grid reference has associated therewith a ray 72 parallel with the z axis and spanning the cubic volume.
 20 Therefore there are $n \times n$ rays, defined in terms of co-ordinates (i, j) , each of which co-ordinates range between 0 and $n-1$.

In that way, a parallel sub-field 70 is defined which

contains a plurality of rays 72 in the vertical direction, which can be used to model light travelling in either direction along each ray in that set of rays.

5 With reference to Figure 3B, the cubic volume 70 previously described is illustrated rotated through an angle ϕ from the vertical about the y axis. No rotation is imparted about the z axis. In that way, each ray of the $n \times n$ array of rays 72 is oriented at an angle ϕ from
10 the vertical. Also, each ray lies in a plane parallel with the x, z plane.

Figure 3C illustrates the cubic volume containing the array of rays further rotated through an angle θ about
15 the z axis. In that way, each of the rays of the $n \times n$ array is oriented at an angle ϕ from the vertical and, when resolved into the x, y plane, is at an angle θ relative the x axis.

20 All possible directions of rays can be described using the following ranges:

$\phi = 0 ; \theta : \text{don't care}$

$0 < \phi < \pi/2 ; 0 \leq \theta < 2\pi$

25 $\phi = \pi/2 ; 0 \leq \theta < \pi$

It should be noted that at $\phi = 0$, the direction of the rays 72 is independent of the value of θ , and that at $\phi = \pi/2$, θ need only pass through a half revolution in order to cover all possible directions.

5

In order to implement the arrangement on a computer, θ and ϕ must be discretised. Two alternative methods of discretising ϕ and θ will now be described.

10 In a first method, R is a predetermined number of subdivisions of 2π , representing the level of resolution which the apparatus is to use. In this embodiment, R is set such that the total number of different directions is around 6000. Then, defining the discretised variables u ,
 15 v corresponding to continuous variables θ , ϕ :

$$\theta = u \frac{2\pi}{R}, \quad u = 0, 1, \dots, R-1$$

$$\phi = v \frac{2\pi}{R}, \quad v = 0, 1, \dots, \frac{R}{4}$$

Therefore, by substituting these equations into the range definitions of ϕ and θ and making u , v the subject
 20 thereof, the ranges of values which u and v can take are as follows:

$v = 0; \quad u : \text{don't care}$
 $v = 1, 2, \dots, (R/4)-1; \quad u = 0, 1, \dots, R-1$
 $v = R/4; \quad u = 0, 1, \dots, (R/2)-1$

5 Since the above formulae refer to the quantity $R/4$, R must be an integer multiple of 4 for u and v to be integers.

10 However, by using the same resolution to divide the range of θ regardless of the value of ϕ , the rays generated by the above scheme for values of ϕ near zero are very close together. In contrast, when ϕ is near $\pi/2$, the generated rays are relatively spaced apart. It would be advantageous to be able to generate a set of rays within
15 the hemisphere defined by the variables $\phi[0, \pi/2]$ and $\theta[0, 2\pi]$, the directions of the rays being equispaced.

20 A second method of discretising θ and ϕ in accordance with a specific embodiment of the invention produces improved spacing of ray directions. In that second method, the discrete values which θ and ϕ can assume are defined as follows:

$\phi = 0, \theta = \text{don't care}$

and

$\phi = (\pi/2R)u ; u = 1, 2, \dots R$

$\theta = (2\pi/3u)v ; v = 0, 1, \dots, 3u - 1$

5 where

$R = 2^{m-1}$

This sampling scheme is governed by a parameter m . The resolution of ϕ and θ increases with the value of m . ϕ is discretised equally over the range $[0, \pi/2]$. θ is sub-divided over the range $[0, 2\pi]$ to a degree which is a variable depending on the value of u , and consequently on the value of ϕ . If ϕ is small, the number of sub-divisions of θ is also small. If ϕ is large, the number of sub-divisions of θ is also large.

For example, if $m = 5$, then $R = 31$. Accordingly, there will be 32 sub-divisions of ϕ over its range. In the lower end of the range of ϕ , for instance if $u = 2$, v will vary between 0 and 5. Therefore, there will be six sub-divisions of θ .

Conversely, at the upper end of the range of u , such as when $u = 30$, v will take values in the range 0 to 89. Accordingly, if $u = 30$, there will be 90 sub-divisions of the range of θ .

This sampling scheme produces sub-divisions of θ and ϕ which provides an approximation to a uniform distribution of rays.

5 Defining rays over the ranges set out above, a complete light field is represented by the co-ordinate system (i , j , u , v). This can be stored in the computer memory 6 as a 4D array, indexed by the co-ordinates i , j , u and v . For efficiency of memory, this 4D array can be further
10 represented in flattened form as a 1D array, being a list based upon ranging through the co-ordinate i , j , u and v .

A scene is capable of being defined within the light field constructed of the arrays of rays. It is
15 convenient to confine the scene for example to a unit cube space defined by the ranges:

- $0.5 \leq x \leq 0.5$
- $0.5 \leq y \leq 0.5$
- 20 - $0.5 \leq z \leq 0.5$

This ensures that any orientation of the cubic volume of the parallel sub-field 70 will fully enclose the scene. That means that no part of the scene is omitted from
25 coverage by rays in any particular direction. In fact, since the cubic volume of the sub-field 70 has smallest

dimension of 2, which is its edge dimension, any scene with longest dimension no greater than 2 can be accommodated within the cubic volume 70. However, at the extreme edge of the light field, the number of rays and ray directions may be limited relative the density thereof in the central region of the light field, so it is preferred to apply a constraint on the size of the scene smaller than the theoretical maximum size of the light field.

10

Assuming that the real ray 48 in Figure 2A matches with a ray in the light field, it is identified by a start point (x_1, y_1, z_1) and a direction (dx, dy, dz) . dx, dy, dz represent small changes in the co-ordinates of a point for a small distance moved in the direction of the ray. This represents the direction of the ray.

(θ, ϕ) or (u, v) are then found using trigonometry on direction (dx, dy, dz) . By applying a rotation matrix corresponding with the identified values of θ and ϕ , the ray is rotated so that it becomes vertical (parallel with the z axis). Then, the (i, j) co-ordinates are found from the point of intersection of this transformed ray with the x, y plane and the previously stated definition of i, j in terms of x, y .

In the event that the real ray 48 does not match with a ray 72 of the light field, it is necessary to find a ray within the light field which best approximates the real ray. It will be appreciated that for any accepted measurement of error and for any accepted error bound, values for m and n , the discretisation variables, can be selected.

Firstly, once the real ray directions has been converted into direction co-ordinates (θ, ϕ) or (u, v) , the direction co-ordinates are "rounded" to the nearest provided value. In the case of (u, v) , this will be the nearest integer. Secondly, once the actual intersection of the real ray, rotated back to the vertical, with the x, y plane has been found, the co-ordinates of the intersection can be "rounded" to the nearest provided value of (x, y) or (i, j) . In the case of (i, j) , this will be the nearest integer. It is preferable to fit the direction first, since errors introduced by approximating direction will have a greater impact on the eventual propagation of light in the scene.

In reverse, a given set of co-ordinates (i, j, u, v) defines a ray 72 which may be within the scene. u and v are used to calculate (θ, ϕ) for that ray 72. Therefore, $(x_i, y_j, -1) R_y(\phi) R_z(\theta)$ and $(x_j, y_i, +1) R_y(\phi) R_z(\theta)$ give the end

points of the ray in real co-ordinates, where R_y and R_z are standard three dimensional rotation matrices about the y and z axes respectively. This is useful for identifying the position of a point lying on the ray
5 within the scene.

The scene 22 of Figures 2A and 2B is illustrated again in Figure 4. A particular plane 74 within the scene is identified by dash lines; the position of the plane 74 is
10 defined by a horizontal line along the left-hand side wall 24 substantially one third of the height of the wall from the top of the wall, and a horizontal line along the right-hand side wall 26 substantially at the foot of the right-hand side wall. The plane 74 intersects the lamp
15 40 and the back of the chair 42. The selected plane 74 contains a large number of rays, depending on the resolution of the discretisation of ϕ and θ . However, Figure 5 shows a selection of these rays 72 in three different directions. All other rays have been omitted
20 for reasons of clarity. The portion of the scene intersected by the plane 74 is shown, with the window 36, the mirror 38, the lamp 40 and the back of the chair 42 illustrated.

25 One particular ray 86 is shown, which passes along the plane through the lamp 40 and the back of the chair 42.

This is illustrated in more detail in Figure 6. The ray, when transformed back to the vertical, extends from $z = -1$ to $z = +1$, with the various objects intersected, namely the left-hand side wall 24, the lamp 40, the back of the chair 42 and the right-hand side wall 26 placed thereon at the appropriate places. All points along the ray can be assigned a position by referring back to the z axis.

Figure 7 illustrates the segment of the ray 86 between the side walls 24, 26 which is the part of the ray which is of most interest, being as it is part of the scene 22.

In order to render an object of the scene into the light field, firstly a ray passing through an object must be computed. Then, when such a ray 86 is computed, all of the intersections of that ray with the object must be identified by position. That position is expressed as a value t in the parametric equation:

$$p(t) = p + t(q-p) \quad 0 \leq t \leq 1$$

where p and q correspond with intersections of the ray 86 with the limits of the scene, i.e. side walls 24, 26. Therefore, as intersections are found, the ray 86 has an associated set of intersections $[0, t_1, t_2, \dots t_k, 1]$

where the t_i are parametric values corresponding to the positions of the intersections with objects. The real co-ordinate position of each intersection can be found from $p(t_i)$.

5

In order to keep account of intersections of the ray, the intersections are arranged in an interval list. Interval lists for all rays of the entire light field are arranged in a four-dimensional array indexed by parameters (i, j, u, v) . Each interval list is initially set to null, before intersections are found and arranged therein. An initial interval list represents the interval $[0, 1]$, i.e. the full length of the ray within the bounds of the scene.

15

Each entry in the interval list includes not only the t value, but other data concerning the nature of the intersection as well. All of the data is collected in a structure called, hereinafter, a T-Intersection, having fields for the t -value, the object intersected, the direction of potential radiance from the intersection, a radiance r , and an unshot radiance u .

20

Unshot radiance is a quantity which is stored at a T-Intersection for an object, and which is used during processing of light propagation in the light field. It

25

represents a quantity of light which has been received by an object intersection, and is to be processed in due course by a further processing step, to be propagated in accordance with the appropriate physical behaviour at that intersection. In a fully processed light field, ready for viewing, unshot radiance is zero at all intersections.

Radiance r and unshot radiance u are typically vectors, expressing energy levels for red, green and blue light (RGB) so that coloured light can be represented. These vectors are initially set to zero vectors. It will be appreciated that other means of expressing coloured light, other than RGB, are also possible.

The direction is set to be either "right" or "left". The "right" direction corresponds with the direction along the ray from $z < 0$ to $z > 0$ when the ray is transformed back to the vertical; "left" is in the opposite direction.

Radiance r and unshot radiance u are vector quantities if the system is arranged to consider colour, polarisation or other characteristics of light within the system which need to be described. The structure of the T-Intersections is illustrated in Figure 8. Advantageously the T-Intersection includes a field for the "normal

vector" at the point of intersection with the object, though this is not essential since the normal can always be computed by reference to the point of intersection and the data defining the object.

5

The T-Intersections should be placed in a data structure which allows for ease of implementation and relatively constant look-up time.

10 In this embodiment, a standard recursive binary tree is used, known hereinafter as an interval tree. The interval tree corresponding to the ray 86 illustrated in Figure 7 is illustrated in Figure 9.

15 In the general case, each node of the interval tree corresponds to a T-Intersection and further comprises two child nodes, called left_Tree and right_Tree. left_Tree and right_Tree are also interval trees. This structure allows left_Tree and right_Tree to have child nodes, and
20 so on. When an object is rendered into the light field, and a first intersection is found with the ray under consideration, the intersection is loaded into the previously empty interval tree for that ray. The interval tree now comprises one node, and the two child
25 nodes are null. As a second T-intersection is loaded into the interval tree, the t-value contained in that

second T-Intersection is compared with the t-value contained in the T-Intersection of the first node. If the t-value of the second T-Intersection is lower than the t-value of the first T-Intersection, then the second
5 T-Intersection is placed in the left-Tree child node of the first node. Otherwise, the second T-Intersection is placed in the right Node child node of the first node. In that way, an interval tree is constructed which is sorted in respect of the t-values, which binary tree can
10 be used with ease to search for T-Intersections by t-value.

A third intersection can be added in the same way - if a child node is full then comparison is made with the t-
15 value of the contents of that child node. Progress is made down the branches of the tree until a null child node at an appropriate position is found.

In the example illustrated in Figure 7, the ray 86 is
20 parameterised such that $t = 0$ at the left-hand wall 24 and $t = 1$ at the right-hand wall 26. The T-Intersections along the ray 86 are identified by subscript, with T_1 to T_4 representing intersections with the chair 42 and T_5 and T_6 representing intersections with the lamp 40. T_7
25 and T_8 represent the intersections with the side walls 24, 26. Accordingly, the T-Intersections have the

following attributes:

$T_1 = (0.50, 42, \text{left}, 0, 0)$
 $T_2 = (0.52, 42, \text{right}, 0, 0)$
5 $T_3 = (0.70, 42, \text{left}, 0, 0)$
 $T_4 = (0.72, 42, \text{right}, 0, 0)$
 $T_5 = (0.07, 40, \text{left}, 0, 0)$
 $T_6 = (0.15, 40, \text{right}, 0, 0)$
 $T_7 = (0.00, 24, \text{right}, 0, 0)$
10 $T_8 = (1.00, 26, \text{left}, 0, 0)$

If those T-Intersections are found in that order, and loaded onto a binary tree, the binary tree will take the form illustrated in Figure 9.

15

Once T-Intersections are loaded into the binary tree, look-up time is dependent on the logarithm of the number of T-Intersections. This is an advantageous arrangement because look-up time in the interval tree will not
20 increase significantly as an increasing number of intersections are loaded thereon. Moreover, as will become apparent from later description, at times it is necessary to find a T-Intersection adjacent a given T-Intersection in a given direction. It is relatively
25 straightforward to find that T-Intersection using the binary tree in its conventional way.

Once all rays have been found which intersect objects, and all intersections have been loaded onto relevant binary trees, the rendering of the scene into the light field is considered complete.

5

Thereafter, radiance must be added to the light field, by taking into consideration any objects within the field which emit light. In the present example, the window 36 and the lamp 40 are considered to be light emitters.

10

Taking into consideration the lamp 40, all of those rays which intersect the lamp are computed. When a ray 86 is computed, a T-Intersection T_5 is identified on the ray as intersecting with the lamp 40. Radiance in accordance with the light emission characteristic of the lamp 40 is added to the data structure, and unshot radiance equal to the change in radiance is also added thereto. T-Intersection T_6 is then found and treated in the same way, and thereafter, all rays intersecting the standard lamp 40 are treated in the same way.

20

Moreover, although not illustrated in Figure 5, rays intersecting the window 36 are treated in the same way with respect to light emitted thereby, or more properly transmitted therethrough. In the interests of simplicity, it is more straightforward to consider the

25

window as an object with light emitting properties than
to render the sun into the scene. However, light
emitting objects could be placed on the other side of the
window, to increase realism, if this is necessary. The
5 embodiment can accommodate transparent objects such as
the glass of the window.

Obviously, if it is desired to represent further objects
on the other side of the window, those objects must be
10 rendered into the light field as well. This may require
some scaling of the scene in order to ensure that the
whole scene is enclosed.

Once all light emission has been added to the light
15 field, each object is considered to establish whether it
is intersected by a ray which carries unshot radiance.
In the example shown in Figures 5 and 7, the back of the
chair 42 is in receipt of unshot radiance along segment
(T_6 , T_1). At that point, the unshot radiance is
20 considered to be emitted light from that point on the
chair, and it is transmitted through diffuse reflection
along all rays emanating from on or near T-Intersection
 T_1 on ray 86.

25 In the same way as a real ray is unlikely to match
exactly a ray in the light field, it is unlikely that

exact coincidence with other rays will occur at intersection T_1 , and so approximations will be necessary. In fact, when a direction for further propagation is identified, the closest ray is selected, and the
5 diffusely reflected radiance and unshot radiance is added to each such ray. Those diffuse reflections are identified by arrows 88, 90, 92, 94 in Figure 5. Moreover, some radiance will be reflected back onto ray 86 towards the lamp 40. Rays 88, 90 are shown despite
10 the fact that they do not coincide with illustrated sets of rays; they will each coincide with another ray not illustrated but within the set of rays in the light field which are contained in the illustrated plane 74. Other diffusely reflected rays will also be identified in other
15 directions not contained in the plane 74.

Once all unshot radiance incident on the chair 42 is dealt with in this way, the next object to be in receipt of unshot radiance is considered. For example, the
20 mirror 38 is now in receipt of unshot radiance as a result of reflected light from the chair along the ray 92. However, in that case the mirror is a specular reflector, and so only one true reflected ray exists. That true reflected ray is illustrated as a broken line
25 denoted by reference numeral 96 in Figure 5. In the example embodiment, no ray within the light field exactly

coincides with true ray 96. Therefore, the parallel sub-field 70 in a direction nearest to the direction of the true ray 96 is identified and the ray 98 in that parallel sub-field 70 closest to the true ray 96 is identified.

5 Then the unshot radiance along the ray 92 is reflected along the best fit reflected ray 98. Therefore, by iterating through all objects in receipt of unshot radiance, radiance can be added at T-Intersections in accordance with the distribution of light in the scene.

10 In fact, the objects marked as being in receipt of unshot radiance are considered in turn, in decreasing order of the amount of unshot radiance received by each object. In that way, objects having most effect on illumination of the scene are dealt with first. The above technique

15 is carried out without regard to any selected viewing position at this point.

The above examples as illustrated in Figures 5 to 9 have been described with regard to solid objects which reflect

20 light either specularly or diffusely, or a combination of both. However, objects which are at least partially transparent can be treated in the same way. An object with transmissive properties will be described in terms of the effect which the object has on light incident

25 thereon. Radiance travelling along a ray intersecting a specularly transmitting surface of an object can be

computed by identifying the point of exit of the light from the object and the true direction thereof, taking account of, for example, the refractive properties of the object. The closest virtual light field ray to this true
5 direction can then be found and radiance can then be propagated along that closest ray.

The incidence of light on a diffusely transmissive object can be computed in the same way, but taking account of
10 the fact that radiance will need to be propagated along a plurality of rays.

A virtual eye 44 is illustrated in Figure 10. The virtual eye 44 comprises a pupil 76, a lens 78, a focal
15 plane 80, an image plane 82 and an optical axis 84. In order to view the scene from this point, all rays identified as entering the pupil 76 are considered to have entered the virtual eye 44. The direction of each ray and its position relative the optical axis 84 are
20 identified, and a lens equation defining the structure and position of the lens 78 is applied thereto. A lens equation is a vector equation which identifies the trajectory of a refracted ray from a given trajectory of an incident ray.

25

Thereafter, the radiance associated with that ray at that

point along the ray is recorded in an array associated with pixels or sub-pixels in order to build up an image. Smoothing functions can be applied to the image so that any parts thereof which are not fully constructed having regard to the number of rays entering the pupil 76 can be filled. It will be appreciated that the step of building up an image is carried out independently of the calculation of the characteristics of the light field with object and light emitters rendered therein, and so that use of the light field is entirely view-independent.

The apparatus of the particular embodiment of the present invention will now be described in more detail with reference to Figures 11, 12 and 13.

15

Figure 11 is a basic block diagram illustrating operational modules of the computer 2 illustrated in Figure 1.

20 A user interface 100 comprises a facility for the interaction of a user with the apparatus. By virtue of the user interface 100, a user can identify the objects which are to be placed in a scene, and their characteristics such as light emission, absorption, colour and/or position etc. Coupled with the user interface 100 is a light field (LF) computation unit 102.

The LF computation unit 102 is operable to define a light field in terms of the four-dimensional co-ordinate system defined above, to render objects into a scene within the light field and to activate light emission such that
5 radiance is distributed through the light field.

A viewer unit 104 is coupled with the LF computation unit 102 so as to obtain data therefrom concerning radiance values along rays 72 of the light field. The viewer unit
10 104 is operative to convert those radiance values into a focused image, and to process that focused image into a form which can be viewed. The viewer unit 104 is linked with the VDU 12 as described with reference to Figure 1.

15 Figure 12 comprises a block diagram showing the components of the LF computation unit 102. The LF computation unit 102 comprises a preordained light field 106 which cannot be modified by the action of the user.

20 In an alternative embodiment, the resolution of the light field 106 could be modified by the user. This could be carried out to select high speed, or high quality imaging depending on user requirements.

25 Via the user interface 100, described with reference to Figure 11, a user defines, or calls up, a predetermined

scene definition file 108. That scene definition file 108 contains information concerning the nature of objects to be placed in the light field 106. Those object definitions are placed in an objects file 110. Object
5 definition information includes surface reflection characteristics, light emission properties, any light transmission properties, and any surface pattern. It also includes a description of the geometry and position of the object in the scene.

10

The LF computation unit 102 further comprises an intersection data computation unit 112 which considers each object in turn from the objects file 110 in respect of the preordained light field 106, to establish the
15 position of intersections along each intersected ray of the light field. The intersection data computation unit 112 is operable to produce an intersection data table 114 for all rays within the preordained light field 106.

20 The LF computation unit 102 also comprises a light computation unit 116 which makes reference to the objects file 110 to identify light emitting objects, and refers to the preordained light field 106 and the intersection data table 114 to render light through the light field.
25 Light is rendered through the light field, as described previously, by up-dating data within the intersection

data table 114.

Figure 13 illustrates in more detail the components of the viewer unit 104. The viewer unit 104 comprises definitions of predetermined viewers 118 (such as a human eye, a regular camera lens, fish eye lens, wide-angle view, zoom etc) which can be selected and/or modified by interface with the user through the user interface 100. The user interface 100 also provides a facility for specifying the position and orientation of the viewer lens to be used within the light field. This facility is a command line interface at which the co-ordinates and viewing direction of the viewer are specified by a user. Alternative arrangements could be provided, such as by means of a pointing device (e.g. mouse, tracker ball) or other orientation device (e.g. wand).

The viewer unit 104 further comprises a viewed image computation unit 120 which makes reference to the selected viewer lens and the intersection data table 114 of the LF computation 102. The viewed image computation unit 120 is operable to produce an image in accordance with a lens equation describing the characteristic of the selected lens. It will be appreciated that a selected viewer 118 could include a sequence of lenses which, in combination, provide a desired optical effect.

Further, the viewer could comprise a light meter capable of delivering data representative of light received at a point, or over a surface. This could be particularly useful in architectural and theatrical design, where
5 light distribution is an important consideration.

Data corresponding to that image is passed to a data conversion component 122, such as for converting the data into raster scan data for use with a VDU. Up to that
10 point, the level of resolution of the image can be defined to suit the resolution of the light field, and need only be converted into a fully pixellated image at the final stage. The converted image is then output to the VDU 12.

15

A plurality of viewers 118 can be defined, either to deliver different images to different VDU's or, in the case of two viewers being provided, to deliver stereoscopic images to screens in a headset for use in
20 Virtual Reality applications.

Figures 14 to 20 describe procedures performed by the LF computation 102 during operation thereof. Figure 14 describes the main procedure of the LF computation 102.

25

On commencement of the procedure, the LF is set up in

step S1-2 by the SET UP THE LF procedure described later. Once the LF is set up, objects are rendered into the LF in step S1-4 by the RENDER OBJECTS INTO LF procedure to be described later.

5

Once all objects have been rendered into the LF, those objects which are light emitters are activated in step S1-6, by means of the ACTIVATE LIGHT EMITTERS IN LF procedure to be described later. Then, in step S1-8, once the light emitters have been activated, radiance is emanated through the LF by means of the COMPUTE REFLECTIONS IN LF procedure described below. Once that procedure has been completed, the LF is fully set up with objects rendered therein and light emission activated.

15 The SET UP LF procedure will now be described with reference to Figure 15. Firstly, in step S3-2, a parallel sub-field 70 is defined as a grid of n by n parallel rays. The rays are arranged in x, y, z space parallel with the z axis and the sub-field is bounded by

20 the following constraints:

$$- 1 \leq x \leq 1$$

$$- 1 \leq y \leq 1$$

$$- 1 \leq z \leq 1$$

25

Following that step, in step S3-4 the sub-field is cycled

through a predetermined number of orientations of ϕ and θ , wherein ϕ is the angle of a ray from the z axis and θ is the angle defined between a ray resolved into the x, y plane and the x axis. In order to cover all possible
5 directions of the rays, ϕ and θ must pass through the following range:

$$\begin{aligned}\phi &= 0 \\ 0 < \phi < \pi/2 ; 0 \leq \theta < 2\pi \\ 10 \quad \phi &= \pi/2 ; 0 \leq \theta < \pi\end{aligned}$$

ϕ , θ are discretised in accordance with the second method identified above, so as to generate a near uniform distribution of ray directions. In step S3-6, a set of
15 rays parallel with the pre-rotated z-axis is generated. The rays are look-ups into a four-dimensional array of interval trees as described above with reference to Figure 9. Step S3-8 enquires as to whether any further combinations of ϕ , θ need to be considered. If so, step
20 S3-10 selects the next combination, and the cycle recommences. If not, the routine returns.

The RENDER OBJECTS INTO LF procedure will now be described with reference to Figure 16. The procedure
25 commences by enquiring in step S5-2 whether any objects are to be rendered into the LF. This is carried out by

calling the OBJECTS file 110. If there are no objects to be rendered into the LF, the procedure returns to the main procedure. Otherwise, in step S5-4, the OBJECTS file 110 is called to obtain the parameters defining the object. Those parameters consist of the shape definition, which is in terms of its geometry, and the surface characteristics of the object. Those surface characteristics consist of the light absorption characteristics, and the bi-directional reflectance distribution function (BRDF) for the surface. The object parameters also include details of any light emission characteristics of the object. At this point, however, only the geometry of the object is taken into account.

Following the obtaining of the parameters defining the object, the procedure calls, in step S5-6, a sub-procedure namely OBTAIN DETAILS OF INTERSECTS OF RAYS WITH OBJECTS 210. When that sub-procedure has completed, the procedure returns to step S5-2 to repeat the enquiry as to whether any objects remain to be rendered into the LF. If no more objects require rendering into the LF, then the procedure returns, otherwise the procedure selects the next object and repeats as necessary.

The OBTAIN DETAILS OF INTERSECTS OF RAYS WITH OBJECT procedure will now be described with reference to Figure

20. Firstly, in step S13-2, the procedure enquires as to whether any rays intersecting the object remain to be considered. If no more rays remain to be considered, the procedure returns to the RENDER OBJECTS INTO LF procedure
5 described above. Otherwise, in step S13-4, an intersected ray is considered and, in step S13-6, intersections of that ray with the object are successively loaded onto a binary tree for that ray. The intersection data consists of the position of the
10 intersection on the ray, the identity of the object intersected, and the potential direction of radiance from the intersection. These data are determined from the parametric equation identified for the rays by means of its co-ordinates (x, y, ϕ, θ) . Further, the intersection
15 data includes radiance and unshot radiance values which are initially set to zero. The procedure then enquires in step S13-8 as to whether any other rays remain to be considered. If no more rays remain to be considered, then the procedure returns to the earlier described
20 RENDER OBJECTS INTO LF procedure. Otherwise, the procedure loops back to step S13-4.

Once all objects have been rendered into the LF, the ACTIVATE LIGHT IMAGES IN LF procedure is called. This
25 procedure is described in more detail with reference to Figure 17. The procedure commences in step S7-2 by

considering an object with light emission properties. Thereafter, a sub-procedure FIND A RAY INTERSECTING OBJECT AND INTERSECTION THEREWITH 212 is called in step S7-4. This procedure will be described later.

5

Once a ray intersecting the object and its intersection therewith has been identified, in step S7-6, an interval is defined which is bounded by that intersection and the next intersection along the ray from that object in the
10 direction of predicted radiance (i.e. "right" or "left") defined in the data associated with the intersection. Then, in step S7-8, the radiance in that interval is set in the data structure associated with the intersection under consideration. That radiance is set in proportion
15 with the light emitting properties of the object, which are called from the parameters defining the OBJECT as held in the objects file 110.

In step S7-10, an enquiry is made as to whether radiance
20 in that interval is greater than a predetermined threshold. This threshold is set having regard to the scaling used for representing radiance, and the threshold for visibility of light in the application for which the light field representation is intended. This threshold
25 for visibility is determined for a graphically represented image as the sensitivity of the human eye,

and for a photometer application to the minimum sensitivity of the photometer.

If the radiance is greater than the predetermined
5 threshold, then the procedure concludes that it would be worthwhile computing any reflections from that radiance. As such, step S7-12 then causes the unshot radiance in the intersection data in question to be set so as to correspond with the change in the radiance. The object
10 at the other end of that interval is identified in step S7-14 and that object is marked in step S7-16 as being in receipt of unshot radiance.

If the radiance in the interval is less than the
15 threshold, then the unshot radiance is not set, and there is no requirement to identify and mark the object at the other end of the interval. Thereafter, any other intersections of that ray with the object in question are considered in steps S7-18 and S7-20 and radiance and
20 unshot radiance are updated as required. Any other rays intersecting the object are then considered in step S7-22, and then in step S7-24 any other objects with light emission properties are considered in the same way. Once all objects having light emission properties have been
25 processed in the same way, the procedure returns.

The ACTIVATE LIGHT EMITTERS IN LF procedure called the FIND RAY INTERSECTING OBJECT AND INTERSECTION THEREWITH procedure. That procedure will now be described with respect to Figure 19. The procedure firstly enquires in
5 step S11-2 as to whether the object is defined in terms of plane polygons. That information is contained in the OBJECTS file 118. If it is so described, a more simplified procedure is followed. A plane polygon is considered in step S11-4, and a ray passing between the
10 vertices of the polygon is found in step S11-6. The ray is marked in step S11-8 as having been found, and the intersection of that ray is found in step S11-10 by Z-interpolation within the polygon under consideration. Once that intersection is found, the procedure returns.

15

The procedure of finding and dealing with rays passing through a plane polygon can be dealt with using a 2-D fill algorithm as set out in "Computer Graphics, Principles and Practice" by Foley, Van Dam, Feiner and
20 Hughes, pp 92-99.

By obtaining intersection positions using a z-interpolation technique, z-buffer data is automatically created which can be used in the creation of an image
25 being a composite of a scene defined in accordance with the embodiment and further objects, which may be dynamic,

which are entered into the scene using a z-buffer algorithm.

In an alternative case, step S11-2 may identify that an
5 object is described not in terms of plane polygons, but
in terms of a geometric shape such as a sphere or an
ellipsoid, or a composite of various primitives.

One method of rendering such an object into the scene
10 involves the definition of a bounding box around the
object in step S11-12, the bounding box being axis
aligned and of smallest size possible while containing
the object. An axis aligned box is convenient for
computation and for finding rays intersecting therewith.

15

A side of the box is considered in step S11-14, and a ray
passing between the vertices of that side of the box is
found in step S11-16. As that ray is found, a check is
made in step S11-18 as to whether the ray has already
20 been marked as found. If so, an enquiry is made in step
S11-20 as to whether any more rays need to be considered.
If not, a new side of the box is considered in step S11-
14. If so, a new ray is found from the same side in step
S11-16. If the enquiry of step S11-18 shows that the ray
25 is not marked as found, then in step S11-24 an enquiry is
made as to whether the ray passes through the object.

This check is required because a ray could pass through the box and not intersect the object. If it does not pass through the object, then the procedure continues from step S11-20, and the same checks are carried out as
5 before.

Once a ray has been found which passes through the object, then the intersections of the ray with the object are calculated with respect to the object shape
10 definition. Once those intersections has been identified, then the procedure returns.

The intersections between rays and the object can be identified in an alternative manner. For every
15 direction, defined by available combinations of ϕ and θ , a direction aligned 2D projection of a bounding box around the object is found. Then, all rays passing through that 2D projection in the direction of ϕ , θ are tested against the object for intersection (if any) and
20 the point along the ray at which intersection takes place.

Finally, once all light emitters in the LF have been activated, and radiance has been added to all of the rays
25 intersecting those light emitters, reflections through the LF are computed. That is carried out by calling the

COMPUTE REFLECTIONS IN LF procedure described now in more detail with reference to Figure 18. Firstly, in step S9-2, the procedure inquires as to whether there are any objects marked as having received unshot radiance.

5 Objects will have been marked as a result of ray intersections having had radiance added to them above a predetermined threshold, and therefore corresponding unshot radiance will have been added thereto as well. If no such objects are found, the procedure returns.

10

An object which has been so marked is put under consideration in step S9-4, and is then unmarked in step S9-6. In step S9-8, the procedure FIND RAY INTERSECTING OBJECT AND INTERSECTION THEREWITH, as previously
15 described, is called. The intersection T_2 of that ray with the object O is then identified in step S9-10. The procedure then searches in step S9-12 through the data structure, which is conveniently in binary tree form, to identify the adjacent intersection T_1 along the ray in
20 the direction of probable radiance, having regard to the BRDF at that intersection, as defined in the intersection data for intersection T_2 . The adjacent intersection T_1 is checked in step S9-14 to establish whether it contains unshot radiance. If it does not contain unshot radiance,
25 then the procedure returns to step S9-8 to consider another ray intersecting the object in question as

before. Once a ray has been identified which has an intersection with the object and which has an adjacent intersection containing unshot radiance, a reflected ray is identified.

5

The step of finding a reflected ray in step S9-16 requires reference to the BRDF of the object, which is contained in the OBJECTS file 110. In this embodiment, the BRDF specifies diffuse reflectance, specular
10 reflectance or a combination of both. In the case of diffuse reflectance, a plurality of reflected rays exist to be found, over the point of intersection of the incident ray under consideration with the object. With a specular reflection, the reflected ray is wholly to be
15 a ray, of the light field, of best fit having regard to a true reflected ray as calculated with respect to the well-known laws of reflection.

In addition to specular and diffuse reflection, in
20 certain embodiments the BRDF can specify more complicated reflection models such as "glossy" reflection, where a cone of reflected rays are produced to correspond with a given incident ray. This effect represents imperfect specular reflection, as described by Phong in CACM 18(6),
25 June 1975, 311-317.

In respect of the or each reflected ray, the intersections of that ray with the object in question are identified in step S9-18. In step S9-20 the change in radiance in accordance with the BRDF is considered as to whether it is greater than the aforesaid predetermined threshold. If it is greater than the predetermined threshold, then, in step S9-22, the radiance in accordance with the BRDF is set in the intersection with the object in question. Then, the aforesaid adjacent object is marked in step S9-24 as having received unshot radiance, and the aforesaid intersection with the object under consideration is updated in step S9-26 so as to reflect the increase in unshot radiance. Otherwise, no consideration is taken of unshot radiance.

15 An enquiry in step S9-28 is then made as to whether any more reflected rays remain to be considered. If so, the procedure returns to step S9-10. Once step S9-10 has been carried out in respect of all reflected rays to be considered, the object is further considered in step S9-30 to establish whether any other rays intersecting that object have received unshot radiance. If so, then the procedure returns to step S9-8. If no more rays are to be considered, then the procedure returns to step S9-2 where the light field is further considered to establish whether any objects remain which are marked as having

received unshot radiance. Once all objects are unmarked, the procedure returns.

Once all unshot radiance from all objects has been
5 considered, the light field is fully defined with respect to the scene rendered therein. At that point, a complete data structure exists independent of any view position or the characteristics of any viewer. It would be possible to use the light field without reference to any viewer,
10 for instance in the design of an art exhibition where even distribution of light throughout a region of the light field needs to be considered. The distribution of light in a gallery could be designed with the above-described arrangement, and the light incident on any
15 particular wall could be monitored by means of a photometer as an alternative to a viewer. In that way, bright patches or dark patches could be eliminated from a particular wall.

20 However the embodiment is also particularly suitable for applications where a particular viewpoint is required, such as is illustrated in Figure 10 with the simulated eye 44. A view can be constructed from the LF through the procedure now to be described with reference to
25 Figure 21. This procedure stands alone from the procedure operating the LF computation, and is controlled

61

by the user interface 100. The viewer procedure commences in step S15-2 by calling for a specification of an aperture of the viewer. This is identified from the VIEWERS file 118 under the control of the user 100. The
5 specification of the aperture includes the actual diameter of the aperture plus its position in space and orientation. Once the aperture has been specified, a set of rays entering the aperture is identified in step S15-4. This is carried out using a procedure such as the
10 FIND RAY INTERSECTING OBJECT AND INTERSECTION THEREWITH procedure previously described.

As each ray is found which enters the aperture, the end point of that ray on the image plane is found in step
15 S15-6. The ray is represented as a vector which is passed, in step S15-6, through a vector equation relating to the composition of the lens. The composition of the lens is described with reference to the VIEWERS file 118. The result of the vector equation is a point on an image
20 plane 82.

The image plane is defined as an array of elements; the radiance associated with the segment of that ray, as called back from the previous intersection of the ray
25 with an object, is mapped, in step S15-8, onto an element within that array. It is possible that several rays will

be incident on any particular element, and as each ray is incident on an element, the radiances are added together.

Eventually, as all rays are considered, an intensity map
5 will be formed on the image plane composed of elements,
which can be processed by filtering and transformation
from the array of elements to form a pixellated image.
In particular, if the lens structure is that of a
relatively simple box camera, the image will need
10 inversion before it can be displayed correctly on a VDU.

The viewing procedure set up above is potentially
relatively fast, and is independent of the level of
complexity of the image. In fact, it is generated in
15 substantially constant time from frame to frame as the
viewer moves. With sufficient and reasonable processing
power, that constant time could be reduced to such an
extent that the viewer could be implemented as a real-
time viewer.

20

The above identified arrangement can also be used in a
headset for virtual reality applications, with
stereoscopic images being projected to the eyes. Both
images can be computed rapidly. Moreover, the
25 orientation of the viewing direction from which images
are computed can be altered depending on the angle of

gaze of the eyes of the wearer of the headset. This is an important development which has previously required significant computational power for its achievement.

5 Although the above apparatus has the potential for a view of a static scene to be changed in real-time, dynamic changes to the scene itself in real-time may be rather more difficult to achieve. However, the apparatus can be combined with other apparatus to define an image of a static, background scene in accordance with the present disclosure, while the other apparatus produces an image which can be changed dynamically, which image is superimposed over the image of the static scene. This might result in a degradation of the realism of the overall image, since there is not likely to be full interaction between the light field and the dynamic objects. As noted earlier, the apparatus can readily be used to retrieve z-values for combination with a z-buffer method for rendering objects into a scene.

20

Moreover, the embodiment described above assumes that the light transmitting medium is non-participatory. In practice, media are rarely non-participatory, and so attenuating properties can be represented by means of placing intersections at random along the rays in the light field. The concentration of those random

25

intersections depends on the level to which the medium in question attenuates light transmitted through it. In that way, the scene can be illuminated taking account of the attenuating properties of the medium. In fact, attenuation is commonly the result of particles (such as water droplets or dust) in suspension in air, and so the random distribution of instructions is a reasonable approximation of reality.

10 The preferred embodiment of the invention as described above has been described in relation to the modelling of a scene to which no changes are to be made following modelling. However, circumstances could arise in which it would be desirable to be able to delete an object from the scene, or to add an object to the scene. Figures 22 and 23 describe procedures by which these two actions can be achieved.

Firstly, Figure 22 illustrates a flowchart defining a procedure for deleting an object from a scene defined in the light field to which light has already been applied. In step S17-2, a ray is identified which intersects the object to be deleted. Thereafter, in step S17-4, a T-Intersection is found for the identified ray carrying radiance incident on the object in question.

Then, in step S17-6, the radiance R along the interval, of which the found T-intersection forms one bound, is noted. The surface normal at the intersection on the object to be deleted is found, and the redistribution of radiance along reflected rays is determined by adding negative radiance (and unshot radiance) to T-intersections for the or each reflected ray. Thereafter, the negative unshot radiance is dealt with as would positive unshot radiance in accordance with the procedure as previously described with reference to Figure 18.

Then, in step S17-8, the radiance at the T-intersection corresponding with the transmitted ray is set to be equal the incident radiance R .

Then, an enquiry is made in step S17-10 as to whether any more T-intersections are to be considered in respect of the object to be deleted. If there are, then the procedure repeats in respect of the next T-intersection.

Then, an enquiry is made in step S17-12 as to whether any further rays are to be considered. If so, the above procedure is repeated in respect of further rays; otherwise, in step S17-16, the intersections corresponding to the deleted object are deleted from the

66

interval trees relating to those rays. Then the previously described routine COMPUTE REFLECTIONS IN LF can be called, in relation to objects marked as being in receipt of unshot energy.

5

Figure 23 is a flow diagram showing steps of the procedure designed to allow an object to be added to a scene rendered into a light field, after light emitters have been activated.

10

Firstly, in step S19-2, the object is located in the field, and a ray intersecting the object to be added is found. Then, in step S19-4, an interval along the ray is identified, bounded by a so-called second intersection with the object to be added, the direction of potential radiance from that second intersection being to the left. The radiance R currently along that interval is looked up in step S19-6, and the surface normal at the second intersection with the object is found in step S19-8.

15

Rays along which that radiance R is to be propagated are determined in step S19-10, from the BRDF for the object at the point of intersection, and in step S19-12 radiance and unshot radiance are added to the intersections of those rays with the object in accordance with the BRDF.

20

Any objects in receipt of unshot energy are marked in step S19-14.

25

The interval along the ray corresponding to the inside of the object is identified in step S19-16, and in step S19-18 the radiance therealong is set to zero.

5 Finally, an interval along the ray is identified in step S19-20, bounded by a so-called first intersection with the object, the direction of potential radiance from that first intersection being to the right. The unshot radiance along that interval is set to be $-R$ in step
10 S19-22, and the object (if any) at the other end of that interval is marked in step S19-24 to be in receipt of unshot energy. The fact that the unshot energy received by the object is "negative" is not relevant; it is only
15 needs to be processed through the light field.

Then, an enquiry is made in step S19-26 as to whether any further rays are to be considered in relation to the object to be added, and if so, the above steps are
20 repeated. If no rays remain to be considered, the procedure COMPUTE REFLECTIONS IN LF is called in step S19-28 in respect of the objects marked as being in receipt of unshot energy (whether positive or negative), following which the ADD AN OBJECT procedure is completed.

25

The above two procedures are optional features, which can

preferably be called from the user interface 100.

These two procedures can be combined in order to provide a procedure for moving an object. In the combined
5 procedure, the union of the object in its original position and the object in its new position is put under consideration, and the rays intersecting this union are found. This is computationally less expensive than
10 considering the deletion and the addition of the object as separate steps to be applied in turn, since many rays intersected by the object in its first position will also be placed under consideration when dealing with the object in its second position.

15 The embodiment is particularly advantageous in that it is capable of providing a globally illuminated scene which can be amended by adding, deleting or moving an object in the scene without the need to recalculate the illumination of the entire scene.

20

Further alternative means of listing the T-Intersections could be provided. For example, an array could be defined for each ray and the T-Intersections could be loaded into this array by t-value. This has the
25 disadvantage of needing to choose a maximum number of intersections along a ray in advance.

Alternatively, the T-Intersections could be stored in a standard linked list, or a doubly-linked list so that it is particularly easy to insert new elements or to delete elements from the list. The disadvantage of this is that
5 there would be linear search time for any particular element, compared to the logarithmic search time required for the binary tree method.

As another alternative, it would be possible to divide
10 the ray into N equal segments. If N is large (say 1000) then any t value is approximated by an index into a particular segment of the ray. Hence, a one dimensional array of T-Intersections, indexed by the index number of the segment corresponding to the intersection, could be
15 used. This has the advantage of constant look-up time (so it is faster than the binary tree method for large number of intersections) but the disadvantage of approximation and memory requirements (since most of the entries in the array would be null). For example,
20 suppose that N = 1000, but that there are only twenty T-Intersections along a given ray. In that case, 980 of the possible entries in the array for that ray would be empty.

25 The memory disadvantages of the alternative set out above can be overcome in at least two different ways. The

first way is to use a run length encoding method. In that case, the storage would be in the form of the number of null entries, followed by the actual non-null entries, and so on throughout the array. However, the look-up
5 time for that array would be non-constant depending on the number and location of non-null entries of the array.

The second way of improving the final alternative set out above is by packing many possible representations into
10 words, rather than considering them as single array entries. For example, if $N = 1024$ (2^{32}), 1024 entries of "0" and "1" can be represented by 32 unsigned integers. Hence, instead of having an array of 1024 potential T-Intersections, an array of 32 unsigned integers, each
15 initialised to zero, is provided. This is less memory intensive. Associated with each such integer is an ordered list of those particular T-Intersections that correspond to "1" entries in the word. So, given a particular t-value, the closest segment in the
20 representation can be found. The quotient of the index of the segment divided by 32 will give a particular array element in which the t-value is located. The remainder of the index of the segment divided by 32 will identify the particular bit that must be set to 1 to represent
25 this t-value. Finally, the T-Intersection can be stored in a linked list associated with this particular integer.

It will be appreciated that the techniques described above are computationally relatively more expensive when considering diffuse reflectors than when considering specular reflectors. That is because diffuse reflection produces a very large number of reflected rays from a single incident ray. If a diffuse reflector is included in a scene and dealt with according to the method described above, the level of computation can become prohibitive. Therefore, the following procedure has been devised with a view to significantly reducing the level of computation required for considering diffuse reflectors within a scene.

For example, Figure 24 shows a perfectly diffusely reflecting object 230, a surface of which is intersected by first and second rays 232, 234 of a light field as described above. The two rays 232, 234 also intersect another body (not shown) which is a light emitter. The two rays 232, 234 intersect a surface of the diffusely reflecting object 230 substantially at the same point.

Therefore, when considering the receipt of unshot energy by the object, and consequent reflection of light, according to the previously described method the first ray 232 is considered and its intersection with the diffuse reflector 230 is identified. Once the

intersection has been found, reflected rays are identified through the hemisphere centred at the intersection and defined by the region bounded by the plane tangential to the surface of the object at the intersection. Radiance and unshot radiance are applied to those reflected rays. Thereafter, the second ray 234 is considered in the same way. However, since the intersection of the second ray 234 with the diffuse reflector is substantially the same as the intersection of the first ray 232 with the diffuse reflector, many, if not all, of the reflected rays in respect of the second incident ray will be the same as those for the first incident ray.

The following method has been devised with a view to reducing the amount of computational effort required in transmitting radiance resultant at a diffusely reflective surface.

Figure 25 shows a schematic diagram of a portion of the surface of the diffuse reflector 230. By way of explanation of the method, the portion of the surface has been divided into a grid comprising a plurality of elements 240. The two incident rays 232, 234 are illustrated as intersecting the object surface within one particular element 240. Once the first ray intersection

has been found, instead of reflecting the ray immediately, the energy associated with that incident ray is accumulated with respect to that element 240. Then, when the second incident ray 234 is calculated as intersecting within the same element, the energy associated with that ray is also accumulated in respect of that element 240. This step results in elements of the surface having accumulated energy associated therewith. The accumulation of energy is carried out according to a weighted sum of terms. Each term represents energy received from a different ray. The weighting of each term is determined with respect to the angle of incidence of the incident ray on the object surface, and the properties of the surface as defined by the BRDF for the surface. In particular, a ray incident or an object at a very acute angle to its surface may be more or less absorbed than one normal to the surface, depending on the nature of the surface.

Figure 26 illustrates the same portion of the object surface, in which a step is illustrated which sends radiance onto rays identified as being in reflection directions from the surface element 240. In this step, the accumulated energy for the surface element 240 is reflected to all identified reflected rays from the surface element 240. In that way, the computationally

expensive reflected energy calculation is carried out only once per surface element 240. Therefore, the level of computation required for calculating diffuse reflection can be reduced.

5

Accumulated energy for an element 240 is stored in an additional field within the structure of the T-intersection associated with that element and relating to a ray of the light field substantially coincident with the surface normal of the object surface at that element of the surface. Practically, it is not possible to add the field only to the set of T-intersections corresponding to surface normals, and so the field is included as part of the structure of the T-intersection data structure. For all T-intersections the accumulated energy field is initialised to zero. For all but the T-intersections corresponding with surface normals, the accumulated energy field will remain set to zero. Therefore, as shown in Figure 27, the T-intersection associated with the ray coincident with the surface normal N at the point of intersection of the incident rays 232, 234 has stored therein the accumulated energy related to the radiance of the incident ray.

25 Thereafter, the accumulated energy is released onto the reflected rays identified in accordance with the BRDF for

the object at that point.

By storing information relating to accumulated energy per element 240 in a part of the T-intersection data structure, one particular advantage of the specific
5 embodiment of the invention is maintained. That is, there is no data held with respect to objects in the scene, only in terms of intersections.

10 The embodiment has several applications, some of which have been identified throughout the description. Other applications will now be described.

Firstly, a model of a building can be rendered into the
15 light field, and illuminated in a desired manner, to verify the quality of an architectural design. The light field data could be conveyed from an architect to a client by a signal, so that the client can review the architect's design. Moreover, lighting can be
20 designed within that rendered building, for use in the production of artistic works, such as exhibitions, stage plays, or film sets.

Secondly, the ability to interchange views and view
25 characteristics enables the use of the apparatus in the design of a sequence of shots for a television or film

production. The invention would be suitable to simulate, in advance of construction thereof, the appearance of a film set under various layout and lighting conditions, and with different types of camera lens.

5

Thirdly, the light field is a digitally encoded representation of a scene, through representation of the light within the scene. A digital encoding is ideal for compression and subsequent transmission, for instance, for broadcast television. A local set top box, hardwired to decode the transmitted digital encoding and capable of extracting images from the light field defined thereby could be used to display games, or other entertainment to people in their homes. The digital encoding could also be transmitted across the Internet.

15

Fourthly, in the case that two images are provided of views of the same light field, those two images can be transmitted to a head mounted display. In that way, stereoscopic images can be formed, which can provide a highly immersive virtual reality environment.

20

In Figure 28, a block diagram is shown of an image processing apparatus having all of the components of the image processing apparatus illustrated in Figure 1.

25

However, in this case, the central processing unit 4 is operable to produce two images of the light field embodied therein. Those two images are produced by two viewers within the light field, having parallel principal
5 optical axes a suitable interpupillary distance apart, which interpupillary distance can be adjusted by a user.

The images produced at those viewers are delivered to the frame buffer 10 from the central processing unit 4.
10 Coupled with the frame buffer 10 is a CD manufacturing suite 250 operable to produce compact discs 252 on which are stored encoded data relating to images delivered from the frame buffer 10. In that way, sequences of images from a light field can be stored on compact disks for
15 viewing later on suitable viewing apparatus.

A head mounted display (HMD) 254 is provided, and is operable to receive both sequences of images from the frame buffer 10. A second VDU 12' is provided, and the
20 two sequences of images can be displayed on those two VDU's 12, 12'.

In this case, the input means 8 can include a manual input device, such as is used commonly in virtual reality
25 applications. In that way, progress of the user through the virtual reality environment can be controlled by the

user through manipulation of the manual input device, and the two images produced by the central processing unit 4 can be adjusted accordingly. Moreover, the interpupillary distance can be measured within the HMD 254, and that measurement can be supplied to the input 8. That interpupillary distance can be translated into the distance between the optical axis of the two viewers within the light field defined in the central processing unit 4 which produce the two images for display in the HMD 254.

The CPU 4 is further coupled to a data compressor 260 which is adapted to receive data relating to the definition of the light field in the CPU 4 and compressing according to a suitable compression algorithm. The data compressor 260 is coupled to a transmitter 262 operable to receive compressed data from the data compressor 260, and to transmit that data by modulation on electromagnetic carrier radiation. The transmitted data is received by a receiver 264 which is in practice implemented as a set top box for a television 266. The receiver 264 comprises a decoder 268, capable of receiving transmitted data and returning it to its uncompressed state, to retrieve the data defining the light field. The data defining the light field is then transferred to a viewer 270 of the receiver 264 which is

of a similar construction to the viewer 104 illustrated in Figure 13 of the drawings. In that way, a definition of a 3D environment can be transmitted to a receiver for manipulation by a user.

5

With reference to Figure 29, the HMD 254 consists of a visor suitable to be placed over a user's head 272. A cable 274 connects the HMD 254 to the image processing apparatus previously described.

10

The HMD 254 conventionally comprises two display units 276 (as illustrated in Figure 30) on which the previously described stereoscopic images can be projected.

15 The described embodiment is particularly suitable for viewing an environment in stereo, since there is little computational effort involved in creating an image from a light field which has already been constructed. Therefore, two views of the same light field can be
20 created without excessive computation. The viewers can be altered in response to monitoring the angle of gaze of the eye of the user, using monitoring means incorporated into the head mounted display 254, to establish objects within the scene which are of interest to the user. In
25 that way, those objects can be brought into focus.

Furthermore, the invention is not limited to the representation of light propagation in a scene. Heat and sound also constitute forms of energy which can be represented as radiating from sources, and so the representation of the propagation of those forms of energy could also be accomplished by the above described technique. In fact, additional fields could be included for heat and/or sound data, in the T-intersection data structure, beyond the existing radiance and unshot radiance fields. In accordance with the invention, heat and unshot heat fields could be added, as could sound and unshot sound fields.

The embodiment can be used for dynamic sources of light, heat and/or sound, since the radiance value at a particular T-intersection can be computed as a function of a nominal source intensity, until the actual source intensity is set. When the actual intensity is set, the actual radiance value at that T-intersection can be ascertained. Therefore, in the case of a field being used for sound propagation determination, propagation may be calculated as a function of a nominal sound intensity, and a received sound intensity can be calculated thereafter by computing the output of that function given a desired sound source intensity.

CLAIMS:

1. Computer apparatus for processing data to model energy propagation within a three dimensional scene, comprising means for defining a three dimensional environment for containing a scene to be represented, the environment comprising a plurality of discrete energy propagation pathways in a plurality of directions within a three dimensional space, means for defining objects and energy sources within said three dimensional space, means for determining intersections between said pathways and said objects and energy sources within said three dimensional space, means for determining propagation of energy along pathways in accordance with said determined intersections, means for defining an energy receiver within said three dimensional space for receiving energy propagated along one or more of said pathways and means for calculating energy received by said energy receiver in accordance with the calculated energy propagation.

2. Apparatus in accordance with claim 1 wherein the energy calculating means is operable to calculate an energy magnitude value on the basis of energy received by the energy receiver.

3. Apparatus in accordance with claim 1 wherein the means for defining an energy receiver is operable to position a viewing plane within said environment, and to determine intersections of pathways of said environment with said viewing plane, said energy calculating means being operable to generate image data on the basis of energy propagation along pathways incident on the viewing plane.

4. Apparatus in accordance with claim 3 wherein the energy calculating means is operable to determine angles of incidence of said pathways with said viewing plane, and to generate said image data in accordance with said angles of incidence.

5. Apparatus in accordance with any preceding claim wherein the means for defining a three dimensional environment comprises means for defining a plurality of subsets of pathways at different orientation within the three dimensional space, the pathways of each subset being parallel.

6. Apparatus in accordance with claim 5 wherein the means for defining subsets is operable to define subsets such that each subset of pathways includes parallel

pathways arranged in a rectangular array.

7. Apparatus in accordance with claim 5 or claim 6 wherein the three dimensional environment defining means includes indexing means, the indexing means being operable to index subsets of pathways in accordance with the direction of the pathways of each subset.

8. Apparatus in accordance with claim 7 wherein the indexing means is operable to index subsets in accordance with spherical coordinates relative a reference plane.

9. Apparatus in accordance with claim 9 wherein the means for defining an environment is operable to define an environment which comprises a larger number of pathways in directions at smaller angles to the reference plane than are defined in directions at larger angles to the reference plane, such that the distribution of pathway directions within the field is substantially uniform.

10. Apparatus in accordance with claim 10 wherein the means for defining an environment is operable to define pathways such that the number of pathways defined in directions at a particular angle to the reference plane

is substantially proportional to the complement of said particular angle.

11. Apparatus in accordance with any preceding claim
5 wherein the means for determining intersections is operable to store information, in respect of an intersection, relating to the identity of the pathway and the object with which it intersects.

10 12. Apparatus in accordance with claim 11 wherein the means for determining intersections is operable to store information, in respect of an intersection, defining energy propagation at that intersection.

15 13. Apparatus in accordance with any preceding claim wherein the energy propagation determining means is operable to process energy propagation information for a pathway with an intersection with an object, to identify one or more pathways on to which energy is to be
20 propagated from said intersecting pathway, and to generate energy propagation information for said identified pathway or pathways.

14. Apparatus for generating data representing energy
25 propagation in a three dimensional scene, the apparatus

comprising means for defining a three dimensional energy propagation environment for containing a scene within which energy propagation is to be represented, the environment comprising a plurality of discrete energy propagation pathways in a plurality of directions in a three dimensional space, along which propagation of energy is to be represented, means for defining objects and energy sources within said three dimensional space, means for determining and storing information defining intersections of said pathways with objects and energy sources within said three dimensional space, and means for determining data defining propagation of energy between pathways in accordance with said determined intersections.

15. Apparatus in accordance with claim 14 wherein the means for defining a three dimensional environment comprises means for defining a plurality of subsets of pathways at different orientations within the three dimensional space, the pathways of each subset being parallel.

16. Apparatus in accordance with claim 15 wherein the means for defining subsets is operable to define subsets such that each subset of pathways includes parallel

pathways arranged in a rectangular array.

17. Apparatus in accordance with claim 15 or claim 16 wherein the three dimensional environment defining means includes indexing means, the indexing means being operable to index subsets of pathways in accordance with the direction of the pathways of each subset.

18. Apparatus in accordance with claim 17 wherein the indexing means is operable to index subsets in accordance with spherical coordinates relative a reference plane.

19. Apparatus in accordance with claim 18 wherein the means for defining an environment is operable to define environment which comprises a larger number of pathways in directions at smaller angles to the reference plane than are defined in directions at larger angles to the reference plane, such that the distribution of pathway directions within the field is substantially uniform.

20. Apparatus in accordance with claim 19 wherein the means for defining an environment is operable to define pathways such that the number of pathways defined in directions at a particular angle to the reference plane is substantially proportional to the complement of said

particular angle.

21. Apparatus in accordance with any one of claims 15 to
20 wherein the means for determining intersections is
operable to store information, in respect of an
intersection, relating to the identity of the pathway and
the object with which it intersects.

22. Apparatus in accordance with claim 21 wherein the
means for determining intersections is operable to store
information, in respect of an intersection, defining
energy propagation at that intersection.

23. Apparatus in accordance with any one of claims 15 to
22 wherein the energy propagation determining means is
operable to process energy propagation information for a
pathway with an intersection with an object, to identify
one or more pathways on to which energy is to be
propagated from said intersecting pathway, and to
generate energy propagation information for said
identified pathway or pathways.

24. Apparatus for analysing energy propagation within a
three dimensional scene, comprising:

means for receiving data defining a three

dimensional energy propagation environment including a plurality of discrete energy propagation pathways in a plurality of directions within a three dimensional space, objects and energy sources within said three dimensional space, intersections of said pathways with said objects and energy sources and propagation of energy along said pathways;

means for defining an energy receiver within said three dimensional space for receiving energy propagated along one or more of said pathways; and

means for calculating the energy received by said energy receiver.

25. Apparatus in accordance with claim 24 wherein the energy calculating means is operable to calculate an energy magnitude value on the basis of energy received by the energy receiver.

26. Apparatus in accordance with claim 24 wherein the means for defining an energy receiver is operable to position a viewing plane within said environment, and to determine intersections of pathways of said environment with said viewing plane, said energy calculating means being operable to generate image data on the basis of energy propagation along pathways incident on the viewing

plane.

27. Apparatus in accordance with claim 26 wherein the energy calculating means is operable to determine angles of incidence of said pathways with said viewing plane, and to generate said image data in accordance with said angles of incidence.

28. A method of processing data to model energy propagation within a three dimensional scene, comprising:
defining a three dimensional environment for containing a scene to be represented;

defining a plurality of discrete energy propagation pathways in a plurality of directions in said environment;

defining objects and energy sources within said source;

determining intersections between said pathways and said objects and light sources;

determining data defining the propagation of energy between pathways in accordance with said intersections;
receiving energy propagation information at a position in said scene for one or more of said pathways;
and

calculating energy received in said receiving step

in accordance with energy propagation data along a selected subset of said pathways.

29. A method in accordance with claim 28 wherein the energy calculating step includes calculating an energy magnitude value on the basis of energy received by the energy receiver.

30. A method in accordance with claim 28 wherein the step of receiving energy includes positioning a viewing plane within said environment, and determining intersections of pathways of said environment with said viewing plane, said energy calculating step including generating image data on the basis of energy propagation along pathways incident on the viewing plane.

31. A method in accordance with claim 30 wherein the energy calculating step including determining angles of incidence of said pathways with said viewing plane, and generating said image data in accordance with said angles of incidence.

32. A method in accordance with any one of claims 28 to 31 wherein the step of defining a three dimensional environment comprises the step of defining a plurality of

subsets of pathways, the pathways of each subset being parallel.

33. A method in accordance with claim 32 wherein the
5 step of defining subsets of pathways includes defining a rectangular array of parallel pathways in each subset.

34. A method in accordance with claim 32 or claim 33
10 wherein the step of defining the three dimensional environment includes indexing the subsets of pathways in accordance with the direction of the pathways of each subset.

35. A method in accordance with claim 34 wherein the
15 step of indexing includes indexing subsets in accordance with spherical coordinates relative a reference plane.

36. A method in accordance with claim 35 wherein the
20 step of defining pathways comprises defining a larger number of pathways in directions at smaller angles to the reference plane than are defined in directions at larger angles to the reference plane, such that the distribution of pathway directions within the field is substantially uniform.

37. A method in accordance with claim 36 wherein the step of defining pathways comprises defining a number of pathways in directions at a particular angle to the reference plane substantially proportional to the complement of said particular angle.

38. A method in accordance with any one of claims 28 to 37 wherein the step of determining intersections includes the step of storing information, in respect of an intersection, relating to the identity of the pathway and the object with which it intersects.

39. A method in accordance with any one of claims 28 to 38 wherein the step of determining energy propagation includes processing energy propagation information for a pathway with an intersection with an object, identifying one or more pathways on to which energy is to be propagated from said intersecting pathway, and generating energy propagation information for said identified pathway or pathways.

40. A method of generating data representing propagation of energy in a three dimensional scene, the method comprising:

defining an energy propagation field for containing

a scene within which energy propagation is to be represented, including defining a plurality of discrete pathways in a plurality of directions along which energy propagation is to be represented;

5 defining objects and energy sources within said source;

 determining and storing information defining intersections of said pathways with objects and energy sources of said scene; and

10 determining data defining propagation of energy along pathways in accordance with said intersections.

41. A method in accordance with claim 40 wherein the step of defining pathways in the three dimensional
15 environment comprises defining a plurality of subsets of pathways, the pathways of each subset being parallel.

42. A method in accordance with claim 41 wherein the step of defining a plurality of subsets includes defining
20 each subset of pathways to include parallel pathways arranged in a rectangular array.

43. A method in accordance with claim 41 or claim 42 wherein the step of defining a three dimensional
25 environment includes indexing subsets of pathways in

accordance with the direction of the pathways of each subset.

44. A method in accordance with claim 43 wherein the
5 indexing step includes indexing subsets in accordance
with spherical coordinates relative a reference plane.

45. A method in accordance with claim 44 wherein the
step of defining pathways comprises defining a larger
10 number of pathways in directions at smaller angles to the
reference plane than are defined in directions at larger
angles to the reference plane, such that the distribution
of pathway directions within the field is substantially
uniform.

15
46. A method in accordance with claim 45 wherein the
step of defining pathways comprises defining a number of
pathways in directions at a particular angle to the
reference plane is substantially proportional to the
20 complement of said particular angle.

47. A method in accordance with any one of claims 40 to
46 wherein the step of determining intersections includes
the step of storing information, in respect of an
25 intersection, relating to the identity of the pathway and

the object with which it intersects.

48. A method in accordance with any one of claims 40 to 47 wherein the step of determining energy propagation includes processing energy propagation information for a pathway with an intersection with an object, identifying one or more pathways on to which energy is to be propagated from said intersecting pathway, and generating energy propagation information for said identified pathway or pathways.

49. A method of analysing energy propagation within a three dimensional scene, including:

receiving and storing information defining a three dimensional energy propagation environment including a plurality of energy propagation pathways in a plurality of directions, including information defining intersections of said pathways with objects and energy sources within said scene and information defining energy propagation along said pathways;

positioning an energy receiver within said environment, and measuring, with said receiver, energy propagation along one or more of said pathways.

50. A method in accordance with claim 49 wherein the

energy calculating step includes calculating an energy magnitude value on the basis of energy received by the energy receiver.

5 51. A method in accordance with claim 49 wherein the step of positioning an energy receiver includes positioning a viewing plane within said environment, and determining intersections of pathways of said environment with said viewing plane, the method further including the
10 step of generating image data on the basis of measured energy propagation along pathways incident on the viewing plane.

15 52. A method in accordance with claim 51 including determining angles of incidence of said pathways with said viewing plane, and generating said image data in accordance with said angles of incidence.

20 53. A computer storage medium storing processor executable instructions operable to configure a computer apparatus to perform the method of at least one of claims 28 to 52.

25 54. A signal carrying processor executable instructions operable to configure a computer apparatus to perform the

method of at least one of claims 28 to 52.

55. A method in accordance with at least one of claims
31, 51 and 52, including the steps of generating a signal
5 conveying image data generated by said image data
generating step and recording the signal either directly
or indirectly.


56. A method in accordance with at least one of claims
10 40 to 48, including the steps of generating a signal
conveying information defining field, pathways,
intersections of said pathways and propagation of energy
along and recording the signal either directly or
indirectly.


15


57. A method in accordance with claim 56, including the
step of performing the method of at least one of claims
49 to 52 on information conveyed on said recorded signal.

1/32




FIG. 1







4/32







FIG. 5



7/32




FIG. 6

8/32




FIG. 7

9/32




FIG. 8

10/32

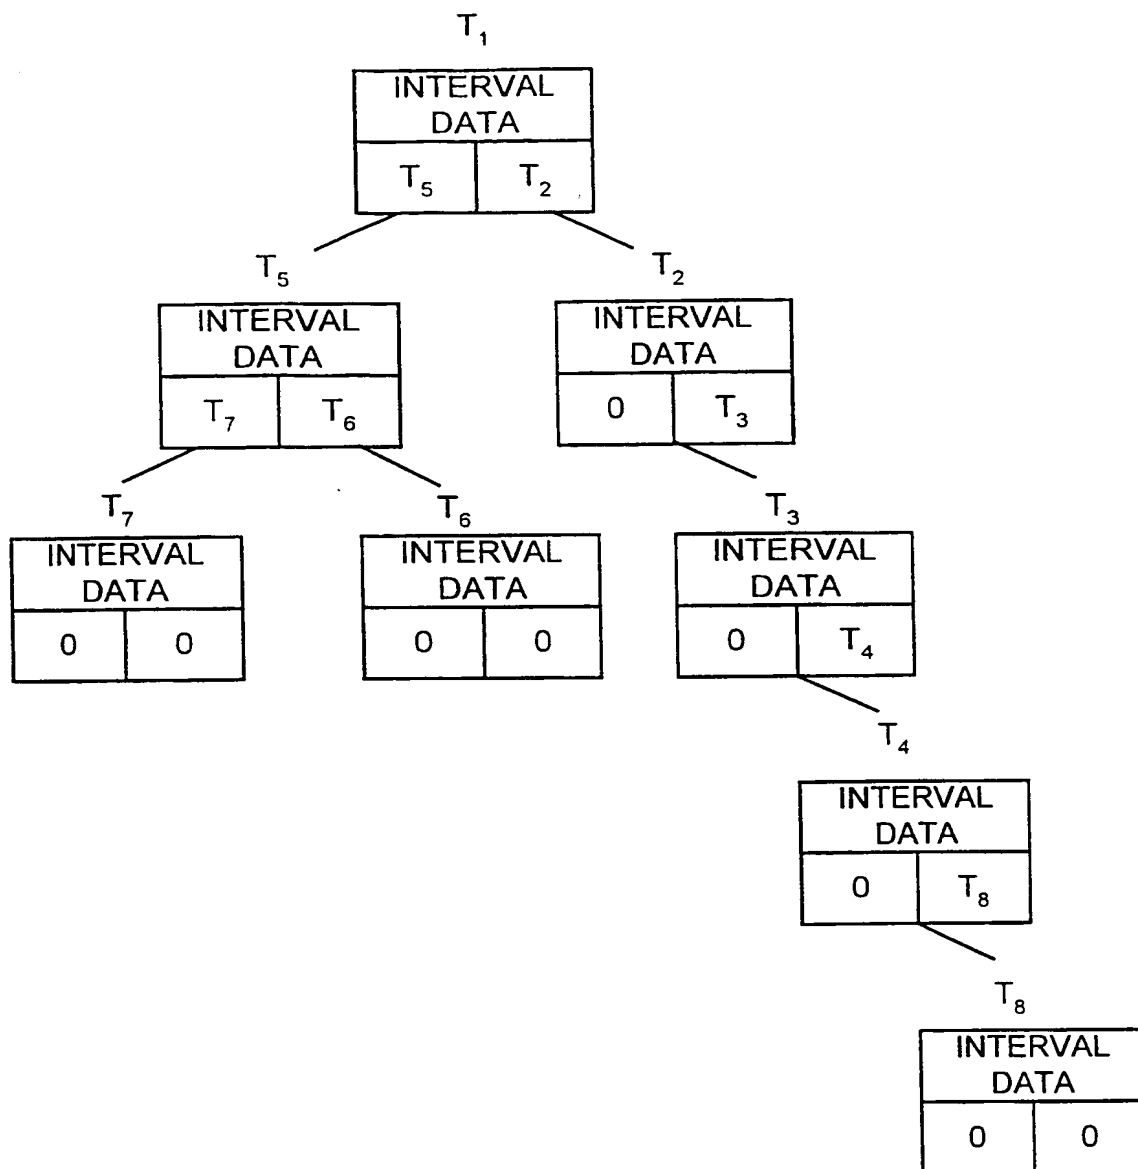


FIG. 9

12/32




FIG. 11

13/32




FIG. 12

14/32

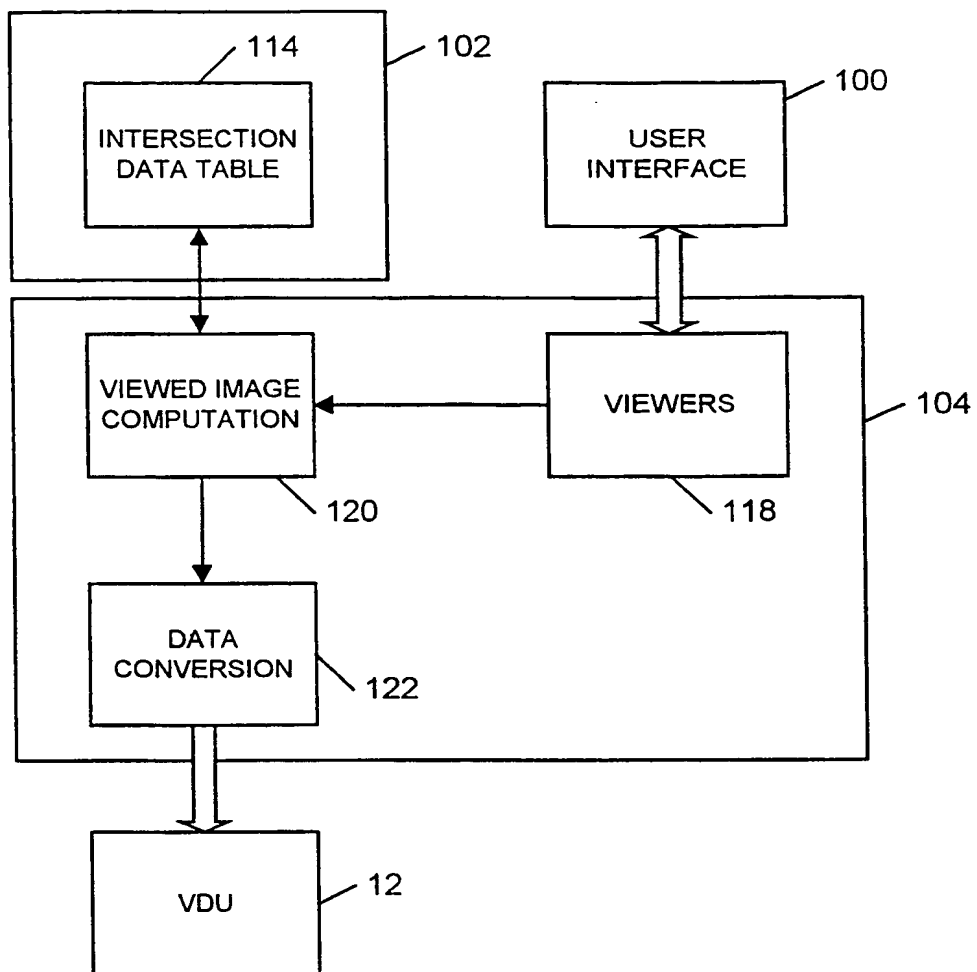


FIG. 13

15/32




FIG. 14

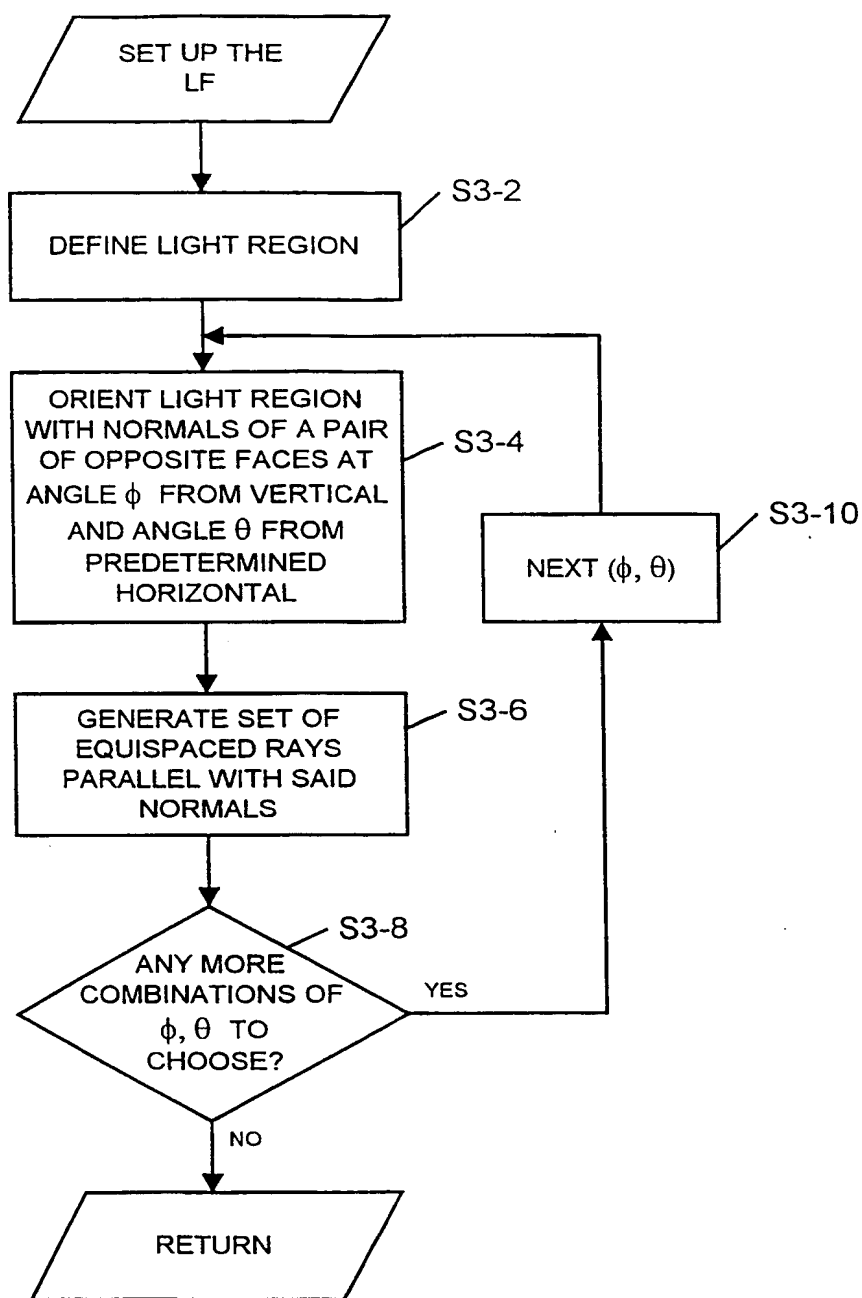


FIG. 15

17/32

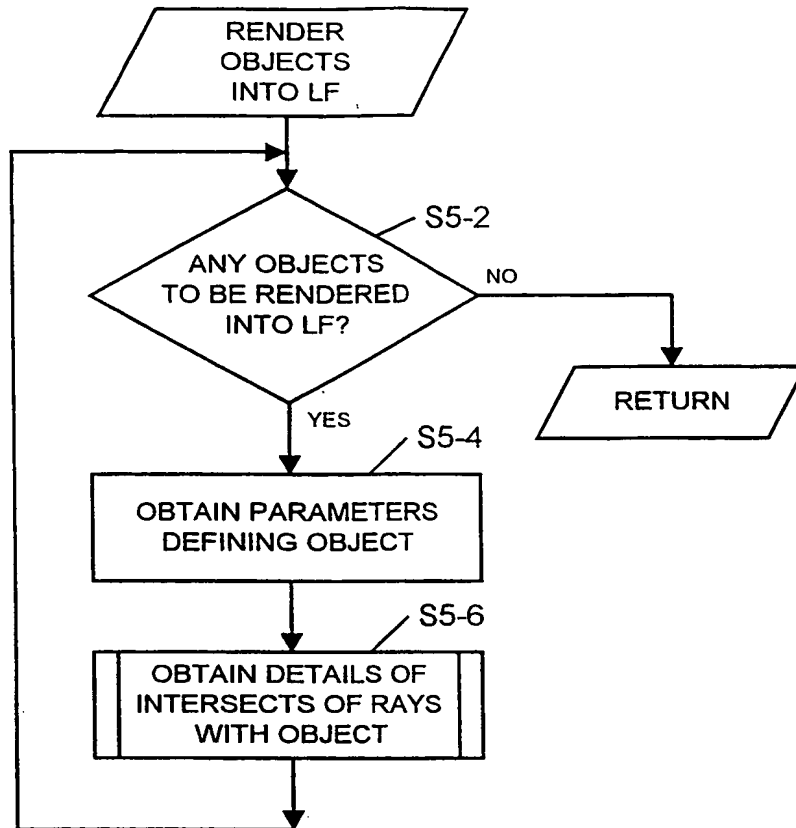
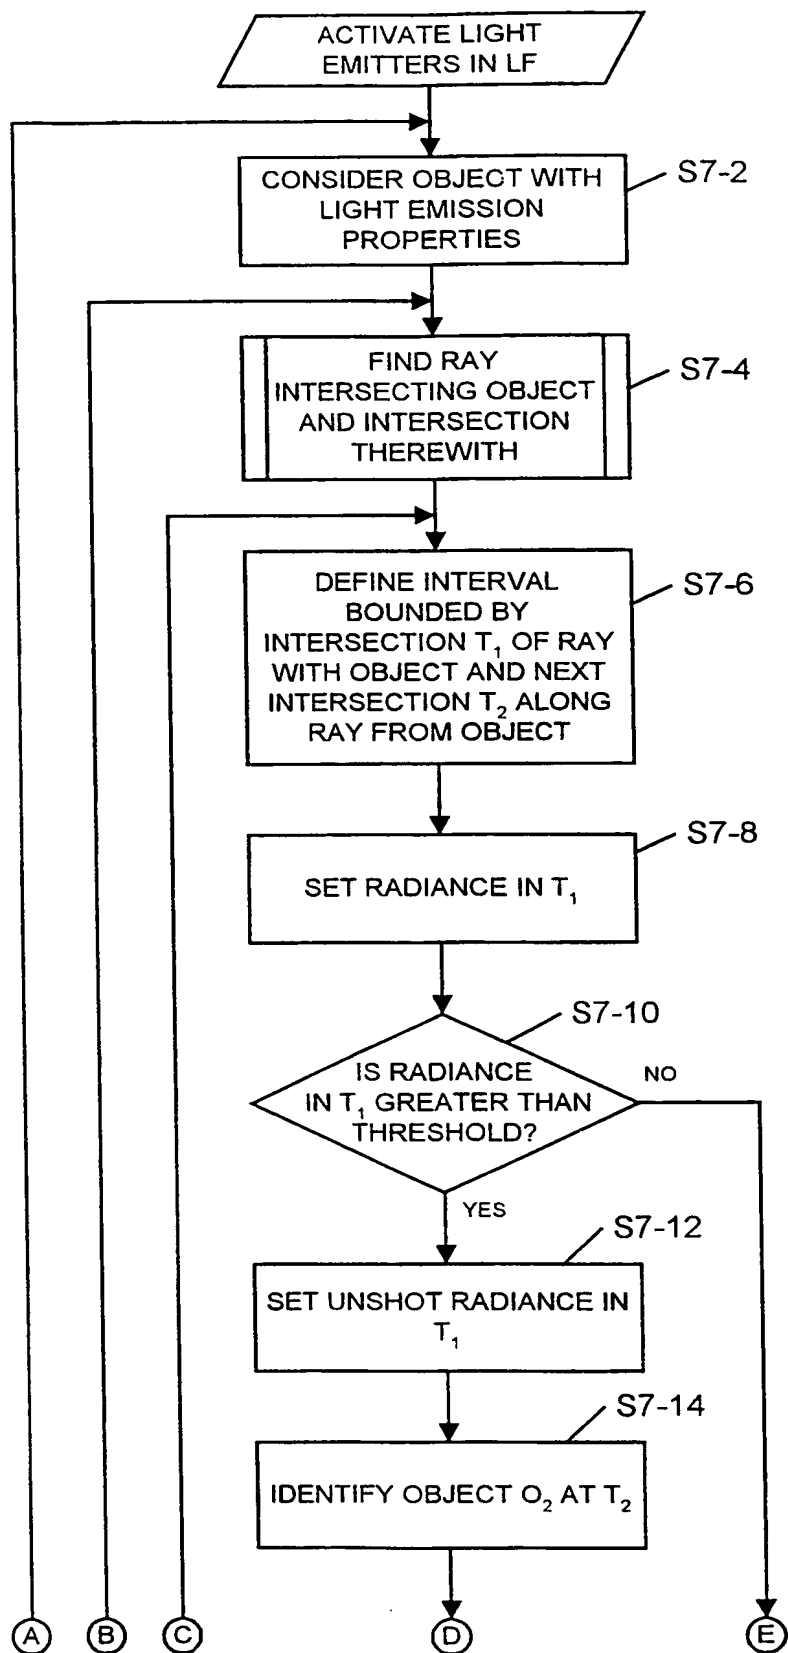


FIG. 16

18/32

FIG. 17



19/32

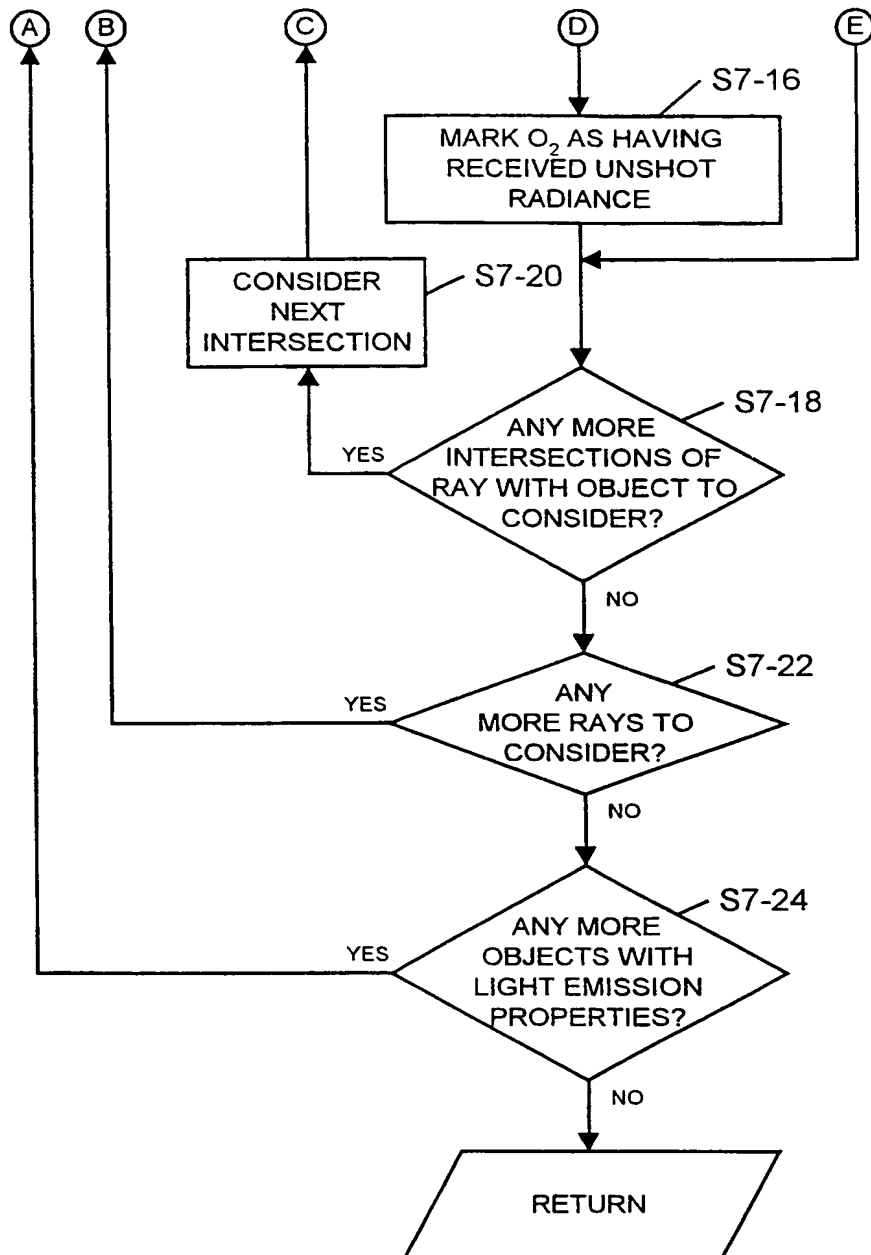
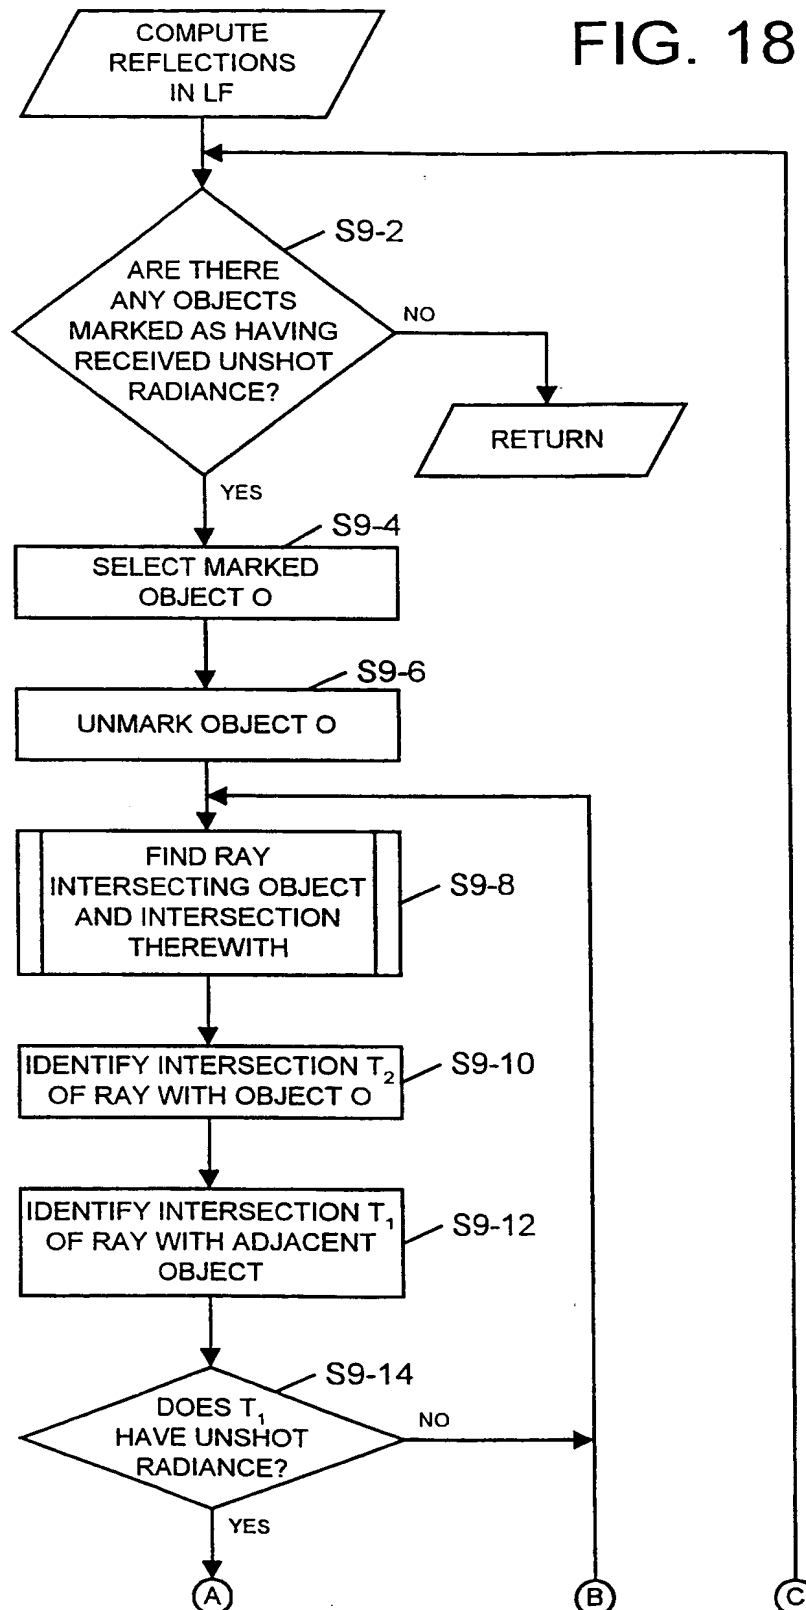


FIG. 17 (cont)

20/32

FIG. 18



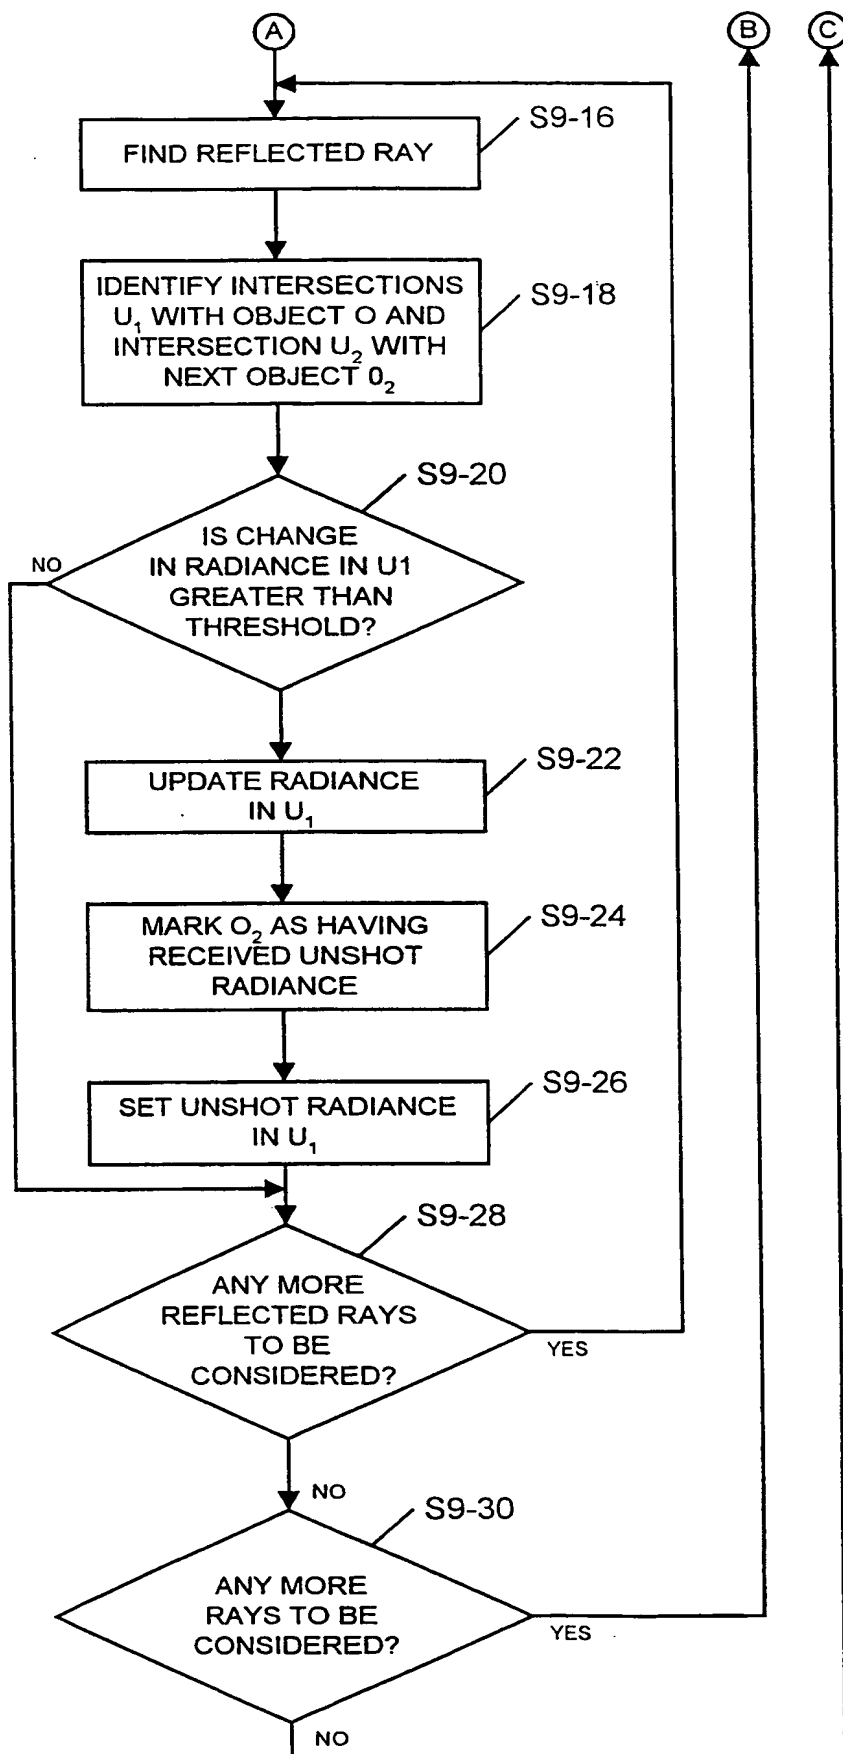


FIG. 18
(cont)

22/32

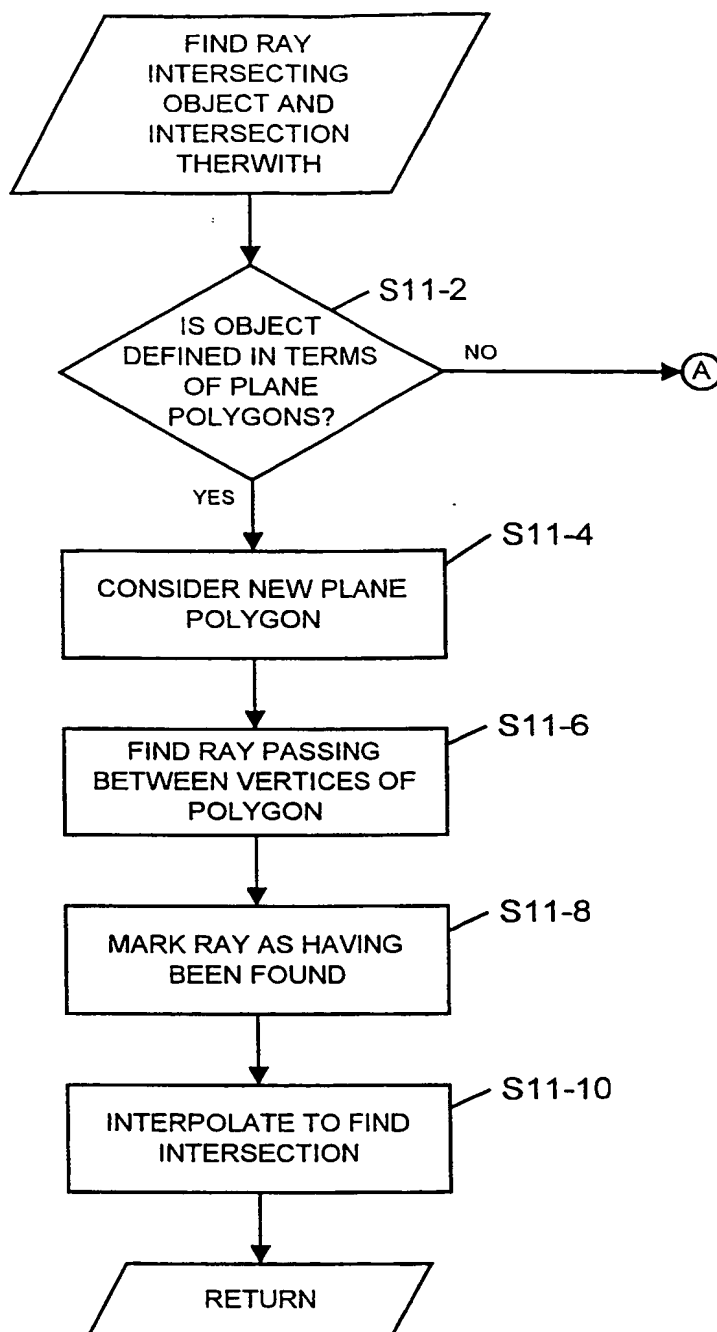


FIG. 19

23/32

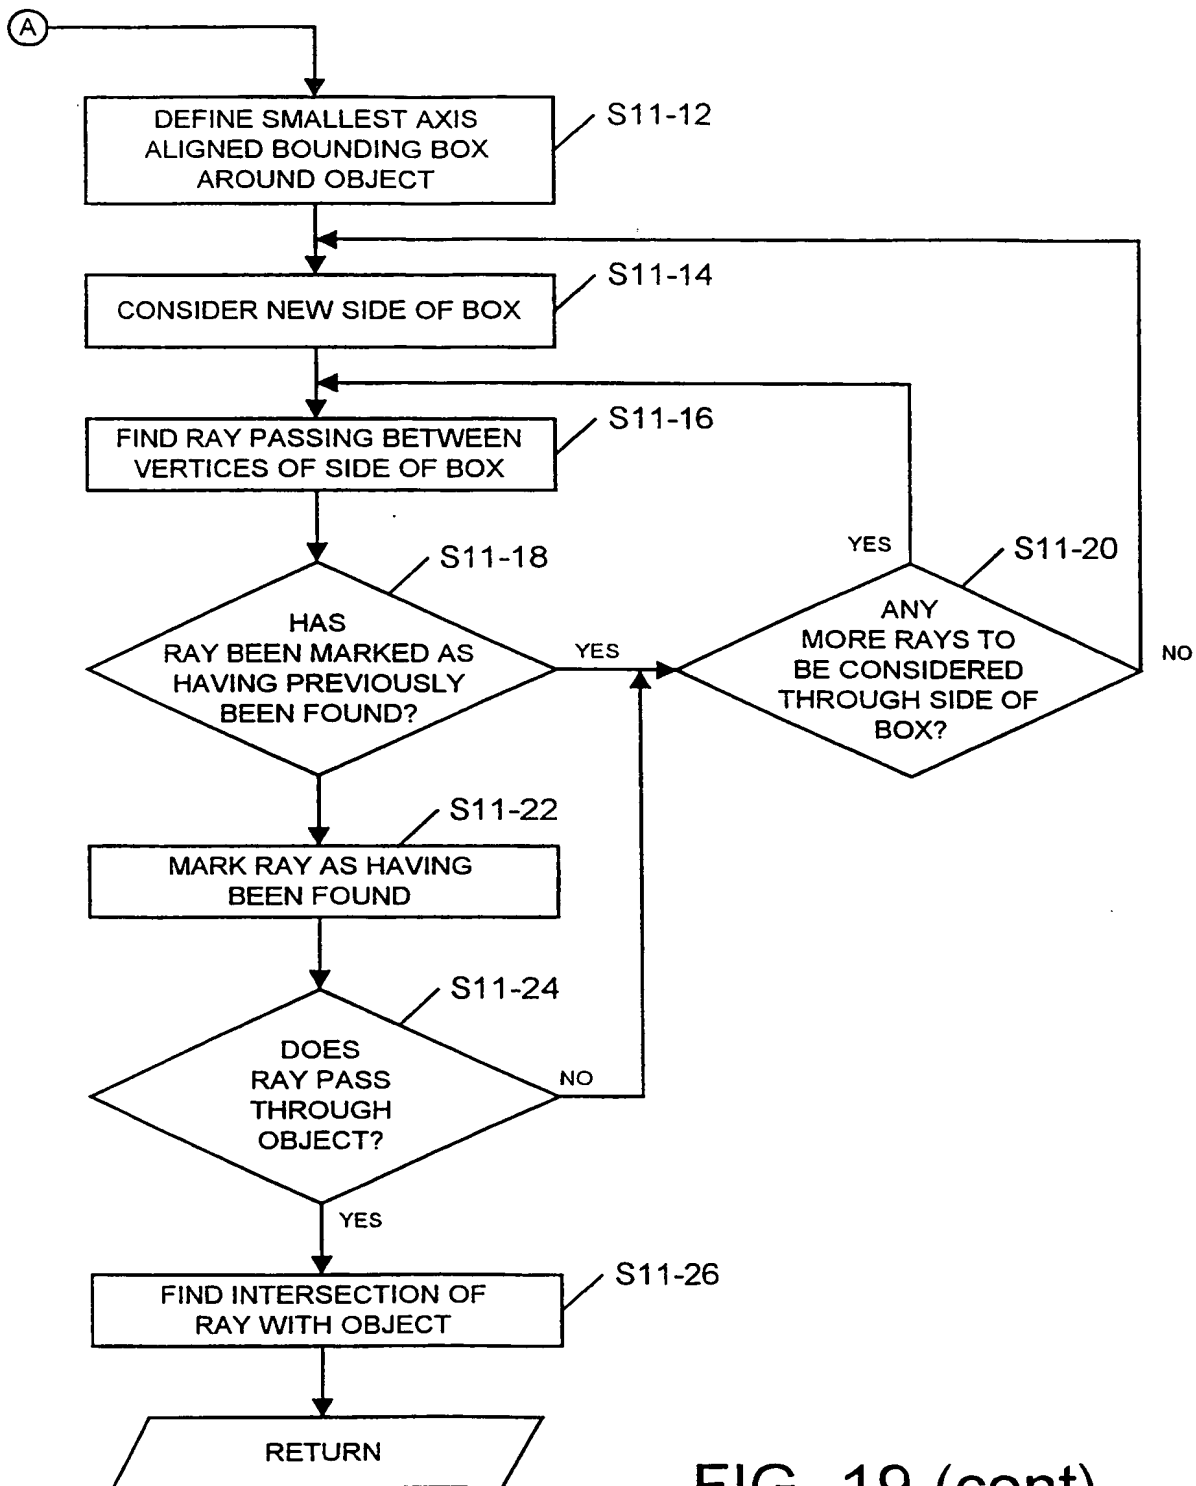


FIG. 19 (cont)

24/32

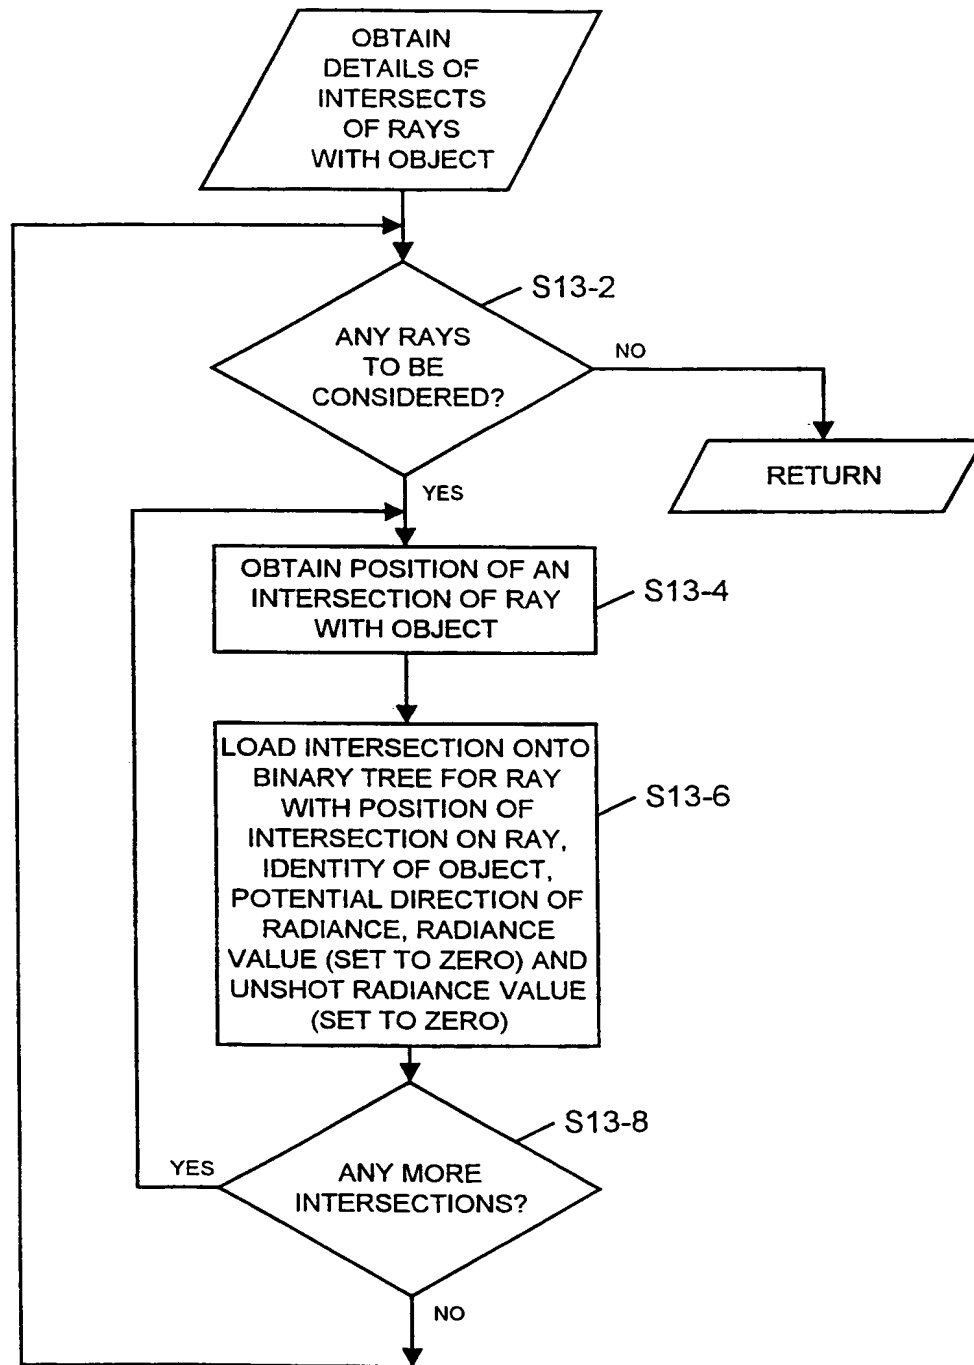


FIG. 20

25/32

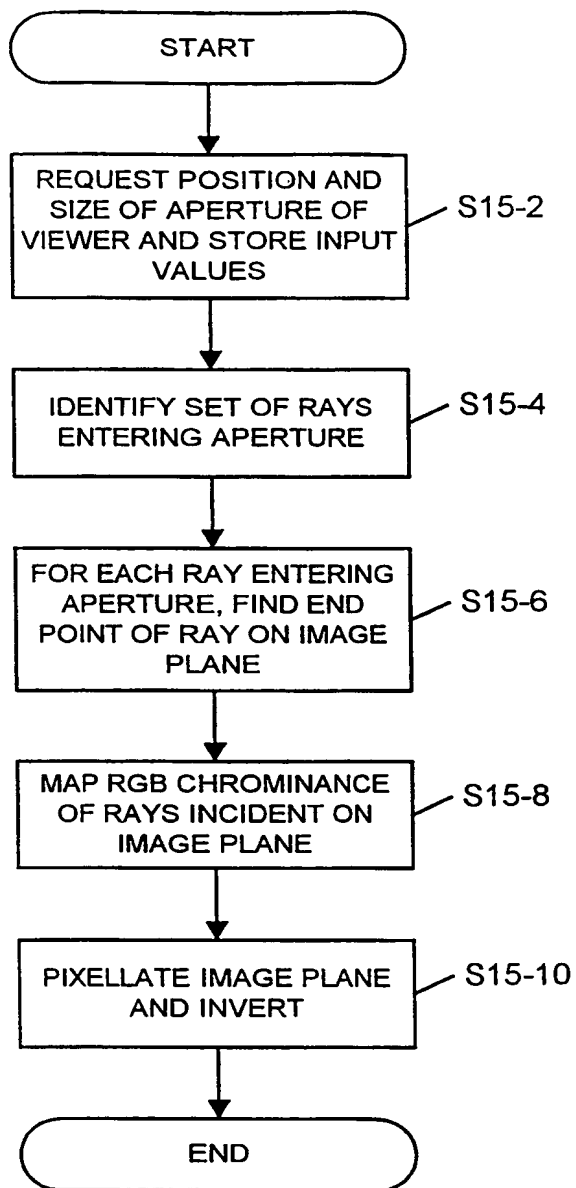


FIG. 21

26/32

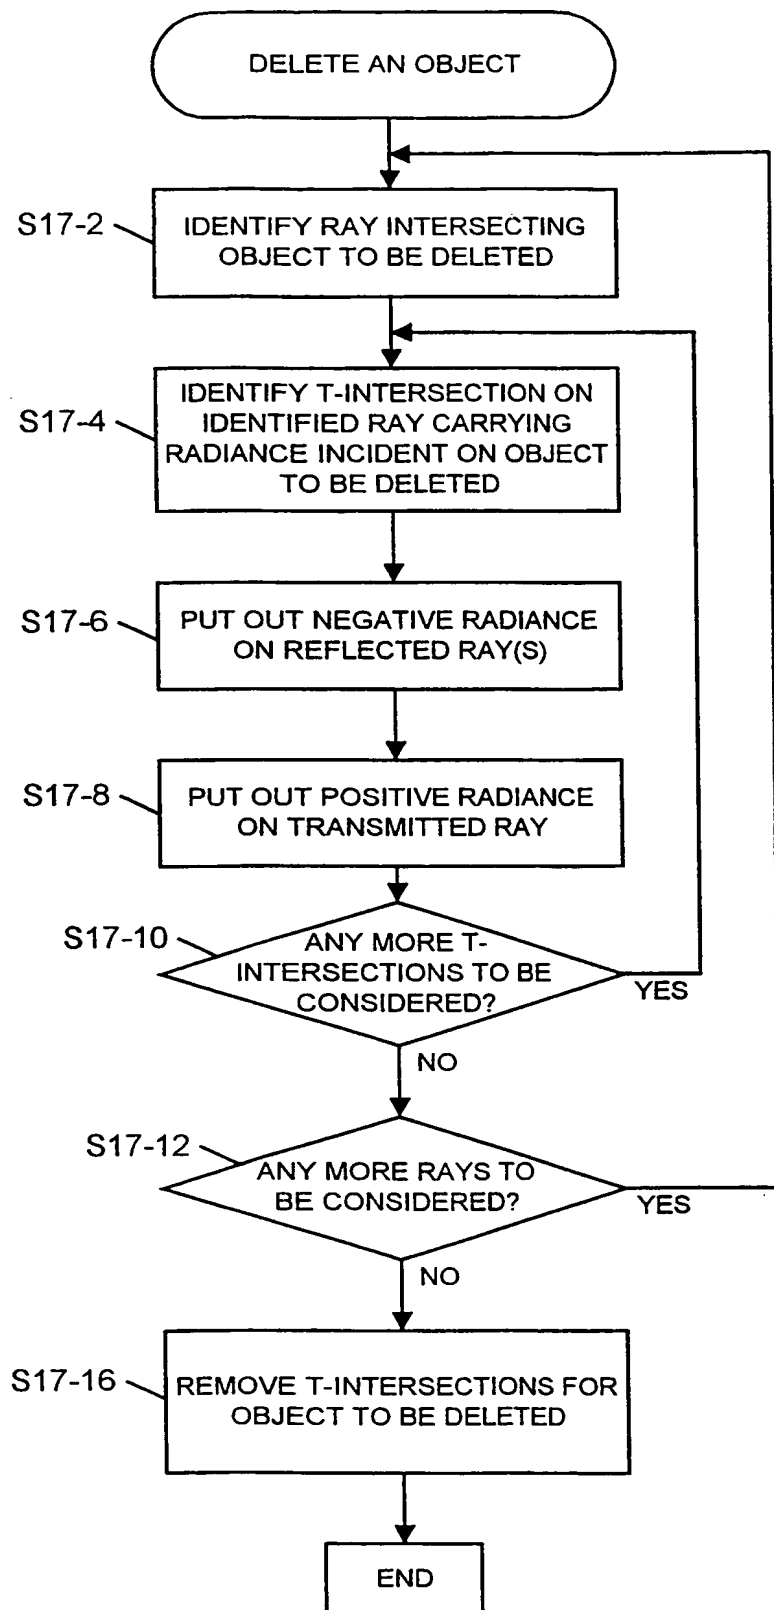


FIG. 22

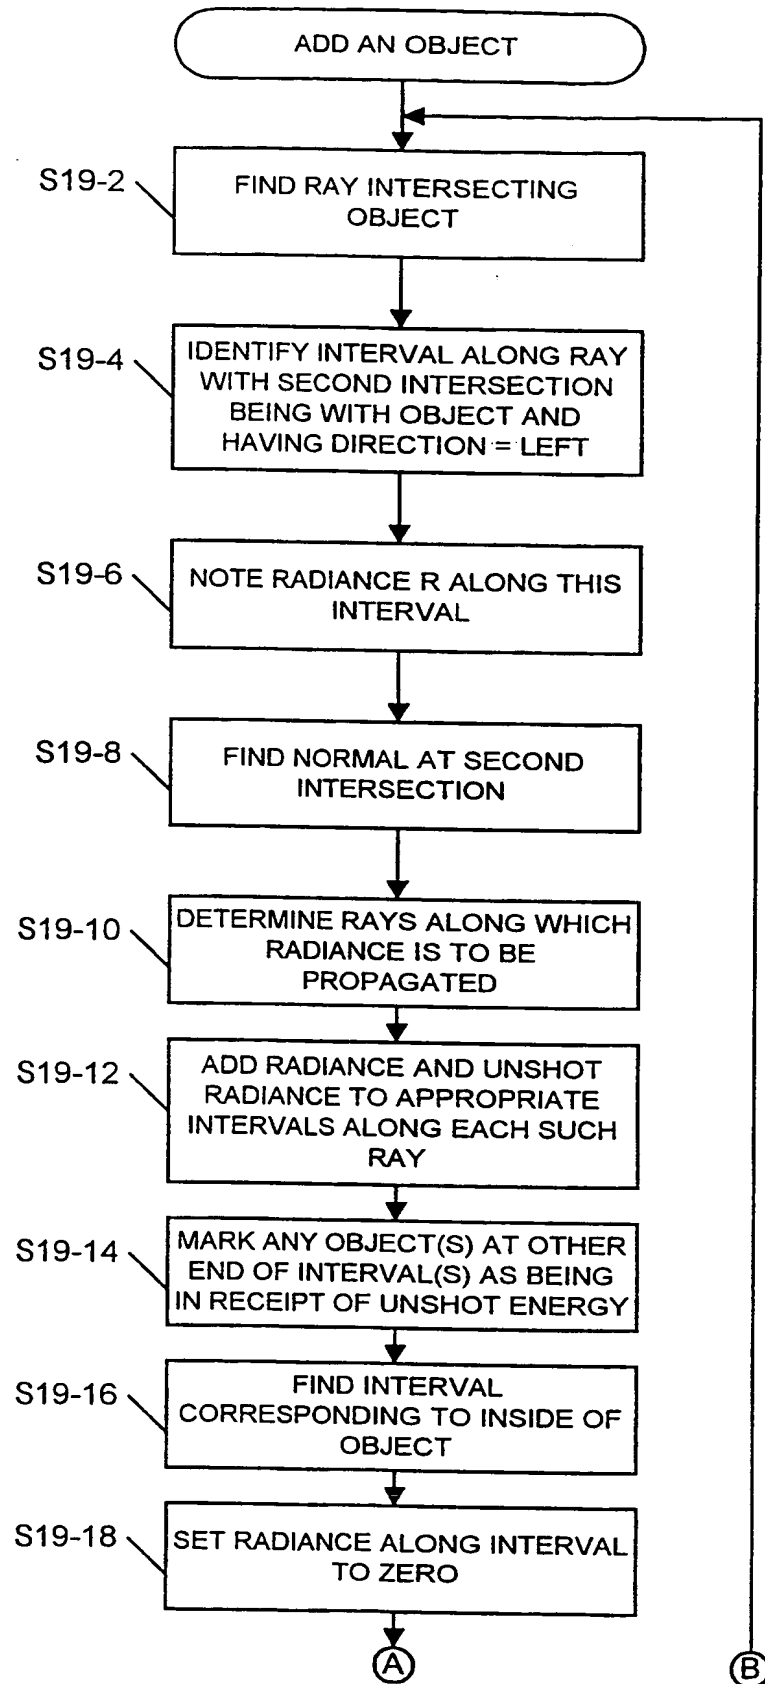


FIG. 23

28/32

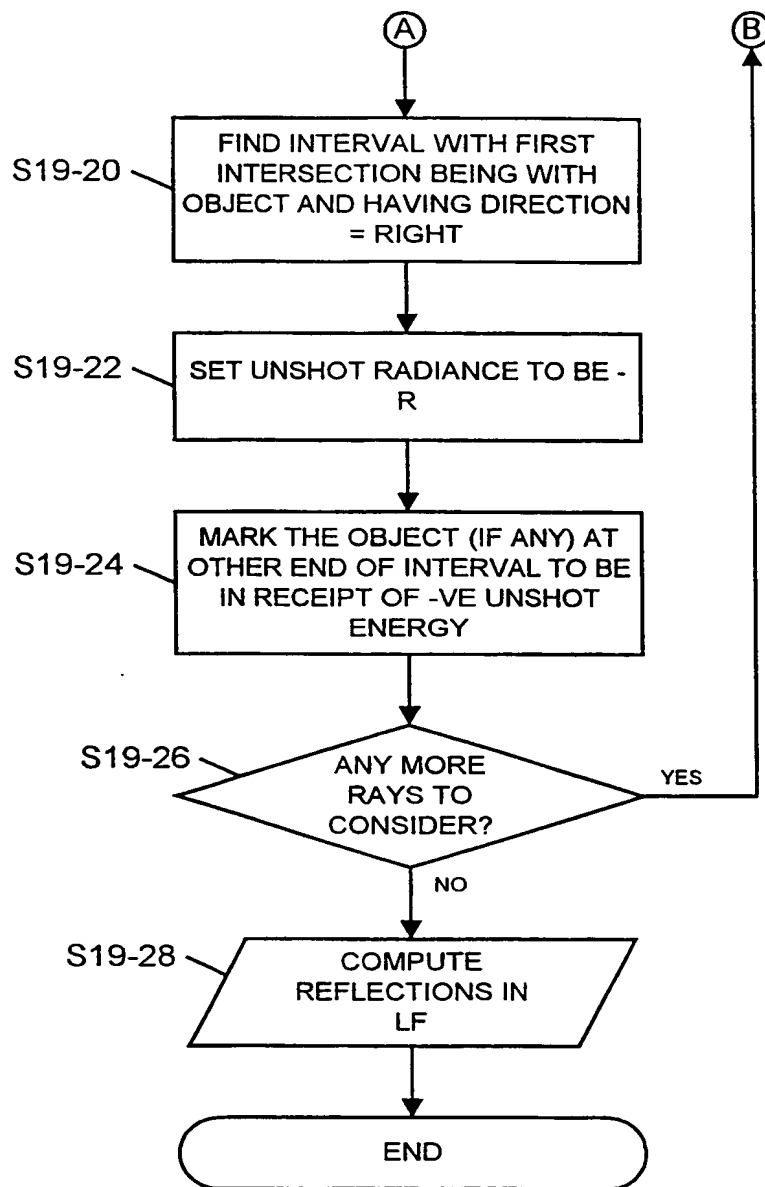


FIG. 23 (cont)

29/32

FIG. 24




FIG. 25




FIG. 26





FIG. 27






FIG. 28




FIG. 29




FIG. 30

INTERNATIONAL SEARCH REPORT

Int. Patent Application No

PCT/GB 99/03229

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06T15/50

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 313 568 A (HAINES ERIC A ET AL) 17 May 1994 (1994-05-17)	1-6, 11-16, 21-33, 38-42, 47-57
Y	abstract; claim 1 column 6, paragraph 1 - paragraph 2 column 10, line 3 - line 8 --- -/--	7-10, 17-20, 34-37, 43-46

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

26 January 2000

Date of mailing of the international search report

03/02/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Fillooy García, E

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 99/03229

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>EP 0 610 004 A (MATSUSHITA ELECTRIC IND CO LTD) 10 August 1994 (1994-08-10)</p> <p>abstract; claim 1 column 3, paragraph 8 -column 5, paragraph 1 column 6, line 49 - line 58</p>	<p>1-6, 11-16, 21-33, 38-42, 47-57</p>
X	<p>US 5 729 672 A (ASHTON MARTIN) 17 March 1998 (1998-03-17)</p> <p>abstract; claim 1 column 1, line 60 -column 2, line 8</p>	<p>1,14,24, 28,40, 49-57</p>
X	<p>EP 0 795 164 B (ADVANCED RENDERING TECHNOLOGY) 9 September 1998 (1998-09-09)</p> <p>claims 1-4 page 2, line 30 -page 3, line 12</p>	<p>1,14,24, 28,40, 49-57</p>
X	<p>US 5 488 700 A (GLASSNER ANDREW) 30 January 1996 (1996-01-30)</p> <p>abstract; claim 1</p>	<p>1,14,24, 28,40, 49-57</p>
Y	<p>L. NEUMANN ET AL.: "Radiosity with Well Distributed Ray Sets" THE INTERNATIONAL JOURNAL OF THE EUROGRAPHICS ASSOCIATION, COMPUTER GRAPHICS FORUM, vol. 16, no. 3, 4 - 8 September 1997, pages C261-C269, XP000870207 Budapest abstract appendix 2</p>	<p>7-10, 17-20, 34-37, 43-46</p>
A	<p>LEVOY M ET AL: "LIGHT FIELD RENDERING" COMPUTER GRAPHICS PROCEEDINGS (SIGGRAPH),US,NEW YORK, NY: ACM, 4 August 1996 (1996-08-04), page 31-42 XP000682719 cited in the application page 31, right-hand column, paragraph 7 page 35, left-hand column, paragraph 1</p>	<p>1-57</p>
	-/--	

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 99/03229

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>NG A: "Assessment of five radiosity acceleration techniques" COMPUTERS AND GRAPHICS, GB, PERGAMON PRESS LTD. OXFORD, vol. 19, no. 5, September 1995 (1995-09) - October 1995 (1995-10), page 727-738 XP004000246 ISSN: 0097-8493 abstract section 3.3 - section 3.5</p>	1-57
A	<p>NEUMANN L ET AL: "RADIOSTY AND HYBRID METHODS" ACM TRANSACTIONS ON GRAPHICS, US, ASSOCIATION FOR COMPUTING MACHINERY, NEW YORK, vol. 14, no. 3, July 1995 (1995-07), page 233-265 XP000554501 ISSN: 0730-0301 abstract section 5.4 - section 5.5</p>	1-57
A	<p>GORTLER S J ET AL: "THE LUMIGRAPH" COMPUTER GRAPHICS PROCEEDINGS (SIGGRAPH), US, NEW YORK, NY: ACM, 4 August 1996 (1996-08-04), page 43-54 XP000682720 cited in the application section 3.1</p>	1-57

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 99/03229

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5313568	A	17-05-1994	EP 0459761 A	04-12-1991
EP 0610004	A	10-08-1994	JP 6223201 A	12-08-1994
			US 5583975 A	10-12-1996
US 5729672	A	17-03-1998	GB 2298111 A	21-08-1996
			GB 2281682 A, B	08-03-1995
			JP 7152926 A	16-06-1995
			US 5596685 A	21-01-1997
			EP 0725367 A	07-08-1996
			GB 2336982 A, B	03-11-1999
			GB 2336983 A, B	03-11-1999
			JP 8255262 A	01-10-1996
EP 0795164	B	17-09-1997	AU 3987795 A	19-06-1996
			DE 69504704 D	15-10-1998
			DE 69504704 T	27-05-1999
			EP 0795164 A	17-09-1997
			GB 2309145 A, B	16-07-1997
			JP 10510074 T	29-09-1998
			US 5933146 A	03-08-1999
			AT 170999 T	15-09-1998
			WO 9617326 A	06-06-1996
US 5488700	A	30-01-1996	JP 7182538 A	21-07-1995